

A Local-Dominance Theory of Voting Equilibria

Reshef Meir^a, Omer Lev^b, Jeffrey S. Rosenschein^b

^a*Harvard University*

^b*Hebrew University of Jerusalem*

Abstract

It is well known that no reasonable voting rule is strategyproof. Moreover, the common Plurality rule is particularly prone to strategic behavior of the voters and empirical studies show that people often vote strategically in practice. Multiple game-theoretic models have been proposed to better understand and predict such behavior and the outcomes it induces. However, these models often make unrealistic assumptions regarding voters' behavior and the information on which they base their vote.

We suggest a new model for strategic voting that takes into account voters' bounded rationality, as well as their limited access to reliable information. We introduce a simple behavioral heuristic based on *local dominance*, where each voter considers a set of possible world states without assigning probabilities to them. This set is constructed based on prospective candidates' scores (e.g., available from an inaccurate poll). In a *voting equilibrium*, all voters vote for candidates not dominated within the set of possible states.

We prove that these voting equilibria exist in the Plurality rule for a broad class of local dominance relations (that is, different ways to decide which states are possible). Furthermore, we show that in an iterative setting where voters may repeatedly change their vote, local dominance-based dynamics quickly converge to an equilibrium if voters start from the truthful state. Weaker convergence guarantees in more general settings are also provided.

Using extensive simulations of strategic voting on generated and real preference profiles, we show that convergence is fast and robust, that emerging equilibria are consistent across various starting conditions, and that they replicate widely known patterns of human voting behavior such as Duverger's law. Further, strategic voting generally improves the quality of the winner compared to truthful voting.

Keywords: Voting equilibrium, Strict uncertainty, Local dominance, Strategic voting

1. Introduction

It is often argued that people vote “strategically”, by trying to promote the election of preferable candidates. Game-theoretic considerations have been applied to the study and design of voting systems for centuries, but the question of how people vote, or should vote, is still open. Suppose that we put aside the complications involved in political voting,¹ and focus a simple scenario that fits all the “standard” assumptions: Each of n voters has complete transitive preferences \prec_i over a fixed set of alternatives M , and each voter’s only purpose is to bring about the election of her most-favorable alternative. We will further restrict ourselves to discussing the common Plurality rule, where the alternative with the maximal number of votes is the winner. This scenario translates naturally to a *game*, in which the actions of each voter are her possible ballots—voting for one of the alternatives, in case of Plurality. One might expect game theory to give us a definitive answer as to what would be the outcome of such a game.

However, an attempt to apply the most fundamental solution concept, Nash equilibrium, to the scenario above, reveals a disappointing fact: Almost *any* profile of actions is a pure Nash equilibrium, and in particular every alternative wins in some equilibrium, even if this alternative is least-preferred by all voters.² This observation triggered a search for more appropriate solution concepts for voting games. These concepts rely on taking into account various additional factors, such as the information available to the voters, collusion and group behavior, and intrinsic preferences towards certain actions. Some solutions focused on variations of the standard single-shot game, for example when voters vote sequentially rather than simultaneously.

Strategic voting. The underlying assumption of game-theoretic analysis is that players are engaged in strategic behavior. But what does it mean to vote “strategically”? Fisher (2004) offers the following definition for what he calls “tactical voting”:³

A tactical voter is someone who votes for a party they believe is more likely to win than their preferred party, to best influence who wins in the constituency.

¹For example, social utilities (Manski, 1993; Brock and Durlauf, 2001), strategic candidates (Calvert, 1985; Feddersen et al., 1990), and other considerations (see e.g. (Riker and Ordeshook, 1968; Edlin et al., 2007)).

²Assuming there are at least three voters.

³Some authors distinct between tactical and strategic voting, especially in political settings. For our purpose they are the same.

The two key components of this definition are *belief* and *influence*. Models of strategic voting differ in how they interpret these terms when considering the behavior of a voter.

Example 1. As a running example, we consider a profile with 3 candidates $M = \{a, b, c\}$. Suppose that there are 100 voters, and that currently votes are divided as: 45 for a , 40 for b , and 15 for c . Among the supporters of c are voters v and v' . Voter v has preference $c \succ b \succ a$, whereas voter v' prefers $c \succ a \succ b$.

While if truthful, both v, v' would stay with c , it seems that c has no chance of winning, and thus a wise strategic decision for v would be to *change* her vote to b . Similarly, v' may prefer to vote for a . A voter's *response function* would dictate what the voter would do in any given state.

Once we define a voter's response function, this immediately induces an *equilibrium model* (or a solution concept): Voting equilibria are simply outcomes where the response of every voter is her current action.

By applying the reasoning above to all supporters of c in the example, we would expect to eventually reach an equilibrium where only a and b get votes. The phenomenon that under the Plurality rule almost all votes divide between two candidates is well known in political science, and is called Duverger's Law (Duverger, 1954).

Our contribution. After enumerating the desiderata we believe should guide the search for a proper solution concept, we review some solutions that have been proposed in the literature, and explain where they fall short of meeting these requirements. We then lay out our epistemic framework, which is a non-probabilistic way of capturing uncertainty. Using this framework and simple behavioral assumptions, we present our response function and equilibrium concept. In the remainder of the paper, we will argue, using formal propositions and empirical analysis, that our solution is indeed the appropriate one for Plurality voting. In particular we show that voters who start from the truthful vote will quickly converge to a pure equilibrium, and that convergence is likely to occur even from arbitrary initial states. We show that in various voter distribution models, using the local-dominance framework enables equilibria with desirable properties — with “better” winners and a “Duverger-like” stable states.

2. Desiderata for Voting Models

We now present some arguably-desirable criteria for a theory of voting. We will not be too picky about what is considered a voting model, and

whether it is described in terms of individual or collective behavior. The key feature of a model is that given a profile of preferences, it can be mapped to a set of outcomes (i.e., of possible or likely voting profiles under the Plurality rule). We classify desirable criteria to the following classes: Theoretic (mainly game-theoretic), behavioral, and scientific.

2.1. Theoretic Criteria

Rationality. The model should assume that voters are behaving in a rational way, in the sense that they are trying to maximize their own utility based on what they know and/or believe.

Equilibrium. The model predicts outcomes that are in equilibrium, for example, a refinement of Nash equilibrium, or of another popular solution concept from the game theory literature. It is more appealing if equilibrium can be naturally computed or even reached by some natural dynamic (similar to the convergence of best-response dynamics to a pure Nash in congestion games).

Discriminative power. The model predicts a small but non-empty set of possible outcomes (sometimes called predictive power). More specifically, it predicts a small set of possible winners, as there may be multiple voting profiles that have the same winner.

Broad scope. The model applies (or can be easily adapted) to various scenarios such as simultaneous, sequential or iterative voting, and to the use of different voting rules.

In addition, we put forward two less formal requirements, that are nevertheless important. First, that predicted outcomes should not include outcomes that are obviously unreasonable or absurd. Second, we would like our model to be grounded in familiar concepts from decision theory, game theory, and voting theory; it will thus be easier to understand, and to compare with other models.

2.2. Behavioral criteria

By behavioral criteria, we mean what implicit or explicit assumptions the model makes on the behavior of voters in the game.

Voters' knowledge. Voters' behavior in the model should not be based on information that they are unlikely to have, or that is hard to obtain.

Voters' capabilities. The decision of the voter should not rely on complex computations, non-trivial probabilistic reasoning, etc.

In addition, we would like the behavioral assumptions, whether implicit or explicit, to be supported by (or at least not to directly contradict) studies in human decision making.

2.3. Scientific criteria

Robustness. We expect the model to give similar predictions even if some voters do not exactly follow their prescribed behavior, if we slightly modify the available information, if we change order of players, etc. Except in a few threshold cases, we would not expect every small perturbation to change the identity of the winner.

Few parameters. If the model has parameters, we would like it to have as few as possible, and we would like each of them to be meaningful (e.g., voters' memory). A related requirement is that the model will be *easy to fit*, given empirical data on voting behavior.

Reproduction. When simulating the model on generated or real preferences, we would like it to reproduce common phenomena such as Duverger's law.

Experiments. The hardest test for a model is to try and predict the behavior of human voters based on their real preferences. By comparing the predicted and real votes (or even just outcomes), we can measure the accuracy of the model.

The behavioral criteria together with the rationality requirement can be thought of as criteria of *bounded rationality*.

Lastly, Some voting models explain how strategic behavior is *better for society*. For example, equilibrium outcomes in Plurality with a particular voter behavior may have a better Borda score, or coincide more with choosing a Condorcet winner. Although this is not exactly a criterion for a good model (a real strategic behavior may not increase welfare), we are interested in the conditions under which the theory predicts an increased welfare, as these may be useful for design purposes.

3. Literature Review

This section is a critical review of some prominent models from the literature of voting under the Plurality rule. For each model, we point out some criteria by which it excels or fails. This is not an exhaustive list, and our purpose is not to criticize other authors, but rather to identify pitfalls of which to beware, and also to identify the positive properties that we would like any new model to preserve. Outside the scope of this review are theories that take into account social utilities, sense of duty, and other incentives that do not follow directly from the voter's preferences (such as the Calculus of Voting (Riker and Ordeshook, 1968)).

The Leader Rule. Before we start, we would like to highlight a model that fairs nicely in almost all of the above criteria, which is Laslier’s *leader rule* for Approval voting (Laslier, 2009). This is a simple parameter-free behavioral strategy, where a voter only needs to take into account a prospective ranking of the candidates (which can be available from a poll, a prior belief, or a previous voting round).⁴ The leader rule is behaviorally plausible, has attractive theoretical properties, makes minimal informational assumptions, and seems to explain human voting behavior at least in some contexts. Unfortunately, it does not seem to have a natural extension to Plurality voting.

3.1. Complete information

The basic notion of a pure Nash equilibrium (PNE) in a normal form voting game is effectively useless, as almost any outcome (even one where all voters vote for their worst candidate) is a PNE. Consider voter v in Example 1. She is powerless to change the outcome, and therefore has no incentive to change her vote. Two refinements that have been suggested rely on plausible behavioral tendencies.

Truth bias. A truth-biased voter gains some negligible additional utility from reporting his true preferences (i.e., his top candidate) (Meir et al., 2010; Dutta and Sen, 2012). He will therefore be truthful, unless he can strictly gain by voting for a different candidate. Dutta and Sen focused on implementation rather than on how truth bias affects a particular voting rule. Nash equilibria under Plurality with truth-biased voters have been studied empirically by Thompson et al. (2013), and analytically by Obraztsova et al. (2013). Indeed, truth-bias significantly reduces the number of pure Nash equilibria (sometimes to zero), and in particular eliminates many unreasonable equilibria such as those where all voters vote for their least-preferred candidate. However if there is gap of at least 2 votes between the winner and the other candidates in the *truthful profile*, then it is a Nash equilibrium with or without truth bias. This is the case in Example 1, and clearly v would still vote for c under truth bias.

Lazy bias. In many voting settings the voting action itself incurs some small cost or inconvenience to the vote. The conclusion that a “rational” voter would often rather abstain (as she is rarely pivotal) is typically referred to in the literature as the “no-vote paradox”, see e.g. (Downs, 1957; Owen and Grofman, 1984). When voting is presented as a normal form game, we can

⁴According to the leader rule, the voter approves all candidates that are preferred to the prospective leader, and approves the leader iff it is preferred to the runnerup.

add abstention as an additional allowed action. A “lazy” voter would thus choose to abstain if she cannot affect the outcome. Pure Nash equilibria with lazy voters were studied, for example, in (Desmedt and Elkind, 2010). These are typically highly degenerated voting profiles, where all voters except one abstain.

There are numerous other models that have been suggested for voting behavior with complete information. Solution concepts vary and include collusion (Sertel and Sanver, 2004; Falik et al., 2012), iterated removal of dominated strategies (Dhillon and Lockwood, 2004), and specific decision diagrams tailored for three candidates (Niemi and Frank, 1982; Felsenthal et al., 1988). Crucially, most of these models assume that (apart from having access to the full preference profile) voters engage in complicated equilibrium computations, perform unlimited steps of iterated reasoning, and so on. That said, models that have been crafted based on empirical observations do manage to replicate interesting phenomena such as Duverger’s law and the election of Condorcet winners (Felsenthal et al., 1988).

3.2. *Voting under uncertainty*

Uncertainty partly solves the the problem of equilibria explosion, since any voter can become pivotal with some probability, and therefore cares about whom to vote for. In Example 1, suppose that v is unsure about the accurate gap between a and b , but knows that it might be zero (whereas the gap between a and c is expected to be much larger). Then the rational thing would be to vote for b , since with positive probability the outcome for v will improve. Uncertainty have also been proposed as a partial solution to the no-vote paradox (Owen and Grofman, 1984).

One model that introduces uncertainty is trembling-hand perfection (Messner and Polborn, 2002), where each vote may be miscounted with some negligible probability. Messner and Polborn also assume some degree of voters’ coordination, and prove that in every equilibrium in their model only two candidates get votes. That is, the model predicts a very strong version of Duverger’s law. They provide some additional results for the three-candidate case.

Myerson and Weber’s (1993) *theory of voting equilibria* provides a different model of uncertainty, one that is closer to our approach. An *outcome* in this model is represented—as in our model—by a prospective vector of candidates’ scores. From this vector they derive a distribution over possible scores, and the outcome is an *equilibrium* if there is a voting profile supporting it, s.t. every voter is maximizing her expected utility w.r.t. this distribution. Myerson and Weber prove that an equilibrium always exists for

a broad class of voting rules including Plurality. Focusing on a few examples with three candidates under Plurality, they show that their model gives reasonable results, and that some equilibria replicate Duverger’s law.

While the Myerson and Weber model is highly attractive in many respects, it suffers from some severe shortcomings. One drawback is that while an equilibrium exists, this is proved by a non-constructive (if elegant) fixed-point argument, and it is not clear how to compute such an equilibrium—let alone how the voters are supposed to find it.

Another restriction is that the model is only defined for non-atomic voters, whose influence on the outcome is infinitesimally small.

However the main issue we find problematic is that voters must engage in non-trivial probabilistic reasoning, even if just to verify that they are playing an equilibrium strategy. The assumption that voters maximize some expected utility function is of course not limited to the Myerson and Weber paper, and is prevalent in the political science literature, see e.g. (Silberman and Durden, 1975; Palfrey and Rosenthal, 1983; Alvarez and Nagler, 2000), as well as in (Messner and Polborn, 2002) which was mentioned above.

From a behavioral perspective such an assumption weakens the model, as people are notoriously bad at estimating probabilities, and are known to employ various heuristics instead (Tversky and Kahneman, 1974).

An additional disadvantage of the expected utility maximization approach, is that voters preferences must be cardinal and cannot be described as a permutation over candidates.

Strict uncertainty. Voting with strict uncertainty (without probabilities) was considered by Ferejohn and Fiorina (1974), who assumed voters each try to minimize their *maximal regret* over all possible outcomes. However their model (like probability-based models) heavily relies on voters having cardinal utilities. Also, they take an extreme approach where voters do not use *any* available information (similarly to the dominance-based approach in (Dhillon and Lockwood, 2004)), and thus all states are considered possible (as also pointed out in a critique on the minimax approach by Aldrich (1993)). Another regret-based model was suggested in (Merrill, 1982).

A different approach to strict uncertainty was taken by Conitzer et al. (2011) who considered manipulations that are weakly helpful in all possible states and strictly helpful in some. Reijngoud and Endriss (2012) looked at a special case where the set of possible states is derived from a poll. Our notion of local dominance is based on similar reasoning, however the mentioned papers restricted their analysis to a single manipulator, whereas we study equilibria. A similar notion of equilibrium under strict uncertainty was defined (with-

out existence or convergence results) by van Ditmarsch et al. (2013). We elaborate on the similarities and differences between these models and ours in Section 7.2.

3.3. Iterative and sequential games

In sequential voting games voters report their preferences one at a time, where every voter can see all of the previous votes (as in Doodle and Facebook polls, and in some internal corporate e-mail surveys). The standard solution concept for sequential games is subgame perfect Nash equilibrium. Subgame perfect voting equilibria, with and without abstention, have been analyzed by multiple researchers (Farquharson, 1969; McKelvey and Niemi, 1978; Dekel and Piccione, 2000; Desmedt and Elkind, 2010). However, subgame perfection is a highly sophisticated behavior that requires a voter to know exactly the preferences of all of her peers. It also requires multiple steps of backward induction, at which human players typically fail (Johnson et al., 2002).

Iterative voting sounds like a similar setting, but has generated a very different type of voting models. In an iterative setting, voters start from some given voting profile, and in each turn one or more voters may change their vote (Meir et al., 2010; Reijngoud and Endriss, 2012). For example, voters can be members of a committee sitting in the same room. The game ends when no voter wants to change her vote. Meir et al. (2010) proved that if voters play one at a time and adopt a myopic best-response strategy they are guaranteed to converge to a Nash equilibrium of the stage game from any initial state. The main problem with this approach is that it does not solve equilibria explosion. In particular, in Example 1 voter v does not have a response that is better than his current action (c). More recent papers on the iterative setting suggested other myopic strategies (Grandi et al., 2013), which suffer from similar problems.

When some or all voters are allowed to change their votes simultaneously, we essentially have repeated polls (Chopra et al., 2004; Reijngoud and Endriss, 2012). A particular model based on uncertainty with iterated polls was suggested by Reyhani et al. (2012). According to this model, each voter considers some set of *possible winners* based on the poll and on some internal parameter called *inertia*, and votes for her most-preferred candidate in this set. However the paper provides few results (focused on three candidates), and applies some arbitrary considerations in the construction of the possible winner set.

4. The Formal Model

Basic notations. We denote $[x] = \{1, 2, \dots, x\}$. The sets of candidates and voters are denoted by M and N , respectively, where $m = |M|, n = |N|$. The Plurality voting rule f allows voters to submit their preferences over the candidates by selecting an action from the set M . Then, f chooses the candidate with the highest score, breaking ties lexicographically.

Let $\pi(M)$ be the set of all orders over M . We denote a preference profile by $\mathbf{Q} \in (\pi(M))^n$. The preferences of voter i are denoted by the total order $Q_i \in \pi(M)$, where $Q_i(a) \in [m]$ is the rank of candidate $a \in M$, and $q_i = Q_i^{-1}(1)$ is the most-preferred candidate. We denote $a \succ_i b$ if $Q_i(a) < Q_i(b)$. Let \overline{Q} be the lexicographic order over candidates. Each voter announces his vote publicly. Thus the action of a voter is $a_i \in M$. The profile of all voters' votes is denoted as $\mathbf{a} \in M^n$, and the profile of all voters except i is denoted by \mathbf{a}_{-i} . When abstention is allowed, we have $a_i \in M \cup \{\perp\}$, where \perp denotes abstaining.

Note that every preference profile \mathbf{Q} induces a game, where the (ordinal) utility of player i in strategy profile \mathbf{a} is $Q_i(f(\mathbf{a}))$. We refer to this game as the *base game*.

If $a_i = q_i$ we say that i is *truthful* in \mathbf{a} , and voter i is called a *core supporter* of a_i . Otherwise, we say that i is a *strategic supporter* of a_i .

The scoring profile $\mathbf{s}_{\mathbf{a}} \in \mathbb{N}^m$ that corresponds to \mathbf{a} assigns a score to every candidate, taking tie-breaking into account. Formally, we refer to $\mathbf{s}_{\mathbf{a}}(c)$ as equal to the number $|\{i \in N : a_i = c\}|$. When comparing two scores, we write $\mathbf{s}_{\mathbf{a}}(c) >_{\overline{Q}} \mathbf{s}_{\mathbf{a}}(c')$ if either $|\{i \in N : a_i = c\}| > |\{i \in N : a_i = c'\}|$, or c, c' have the same number of votes and $c \succ_{\overline{Q}} c'$. We usually omit the subscript from $>_{\overline{Q}}$ as it is clear from the context.

We will use \mathbf{a} and $\mathbf{s}_{\mathbf{a}}$ interchangeably where possible, sometimes omitting the subscript \mathbf{a} (note that we may only use \mathbf{s} in a context where voters' identities are not important). Given a state \mathbf{s} and an additional vote a_i , in the concatenated state $\mathbf{s}' = (\mathbf{s}, a_i)$ we have $s'(a_i) = s(a_i) + 1$, and $s'(a) = s(a)$ for all $a \neq a_i$.

4.1. An intuitive description of voter response

While the notation we will introduce momentarily is somewhat elaborate and is intended to enable rigorous analysis, the main idea is very simple and intuitive. We lean on two key concepts that are featured in previous models: *dominated strategies*, and *best-response*. From the perspective of voter i , a candidate a *dominates* candidate b if $f(\mathbf{s}, a) \succ_i f(\mathbf{s}, b)$ for *all* \mathbf{s} . In contrast, a is a *better-response* for a voter voting for b in a *particular state* \mathbf{s}^* , if $f(\mathbf{s}^*, a) \succ_i f(\mathbf{s}^*, b)$.

In our model, we will relax both concepts in a way that takes into account voters’ uncertainty over the actual outcome. We assume that voters have a common *estimated*, uncertain, view of the current state \mathbf{s} . In any given state, a voter considers a set of multiple “close” states which might be realized **without assigning probabilities to them**. We say that *a locally dominates b* in \mathbf{s} if $f(\mathbf{s}', a) \succ_i f(\mathbf{s}', b)$ in all \mathbf{s}' that are considered “possible” in \mathbf{s} . Our key behavioral assumption is that a voter will avoid voting for candidates that are locally dominated (a standard assumption in strict uncertainty models, see Section 7.2). In an iterative setting, a voter will vote for her most preferred candidate—among those who locally dominate her current action. Our key epistemic assumption (which is new) is that the possible states are those that are “close” to \mathbf{s} according to some reasonable metric over vote counts.

Tying our model back to Fisher’s definition of tactical voting, a voter’s *belief* is captured by the estimated state \mathbf{s} , whereas her *influence* is reflected by local dominance. The different sets of states that voters consider are part of their type, and can account for diverse voter behavior, yet ones that are bounded-rational. Consider Example 1, where the estimated counts are $\mathbf{s} = (45, 40, 15)$. If voter v also considers as possible states where scores vary by ± 10 , he will be better off voting for b . Voting for b might influence the outcome, whereas voting for c is futile (unless v considers an even higher variability in scores).

One assumption that requires justification is the existence of the publicly known, estimated state \mathbf{s} . In an iterative and sequential settings the shared common view is easy to explain, as the estimated state is the actual current voting state. Uncertainty still exists since some voters might change their vote in later rounds. In a simultaneous-vote game, the estimated state might be due to prior acquaintance with the other voters or due to polls, and uncertainty is due to polls’ inaccuracy. While the assumption that voters know the (estimated) current state is slightly stronger than the one made by Laslier’s heuristics for Approval voting (where only the identity of the leader is important), it is much weaker than the assumption that a voter knows the entire preference profile, or the exact distribution over preferences.

4.2. Local dominance

Let $S \subseteq \mathbb{N}^m$ be a set of states.

Definition 1. We say that action a_i *S*-beats a'_i (w.r.t. voter i) if there is at least one state $\mathbf{s} \in S$ s.t. $f(\mathbf{s}, a_i) \succ_i f(\mathbf{s}, a'_i)$. That is where i strictly prefers $f(\mathbf{s}, a_i)$ over $f(\mathbf{s}, a'_i)$.

We can think of S as states that i *believes* to be possible (where these states do not include the action of i himself). The definition, however, does not depend on this interpretation.

Definition 2 (Local dominance). We say that action a_i S -dominates a'_i (w.r.t. voter i) if (I) a_i S -beats a'_i ; and (II) a'_i does *not* S -beat a_i .

Note that S -dominance is a transitive and antisymmetric relation (but not complete). See more on epistemic interpretation in Section 7.2, where we also compare our definition of local dominance with previous work. In particular our definition coincides with the definition of dominance in (Conitzer et al., 2011) and with similar definitions in (Reijngoud and Endriss, 2012; van Ditmarsch et al., 2013). The novelty comes from the way S is constructed, which is explained next.

4.3. Distance-based dominance

Recall that every full or partial profile $\mathbf{a} \in M^n$ corresponds to state $\mathbf{s}_{\mathbf{a}} \in \mathbb{N}^m$. Suppose we have some distance metric for states, denoted by $\delta(\mathbf{s}, \mathbf{s}')$. For voter i and $\mathbf{a} \in M^n$, let $S_i(\mathbf{a}, x) \subseteq \mathbb{N}^m$ be the set of states that are at distance at most x from \mathbf{a}_{-i} . Formally, $S_i(\mathbf{a}, x) = \{\mathbf{s}' : \delta(\mathbf{s}', \mathbf{s}_{\mathbf{a}_{-i}}) \leq x\}$.

The δ distance may be a ℓ_d norm for some $d \geq 1$. Thus $\delta_{\ell_1}(\mathbf{s}', \mathbf{s}) \leq x$ means that we can attain \mathbf{s}' by adding/removing a total of x voters to \mathbf{a}_{-i} (think of x as an integer, and note that the total number of votes in \mathbf{s}, \mathbf{s}' may be different). Similarly, the ℓ_∞ norm means that we can add or remove at most x votes for each candidate.

Another distance we can consider is the *multiplicative distance*, where $\delta_M(\mathbf{s}', \mathbf{s}) \leq x$ if for all $a \in M$, both $s'(a) \leq s(a)(1+x)$ and $s(a) \leq s'(a)(1+x)$. Intuitively, this means that the score of each candidate can change (either increase or decrease) by a factor of at most $(1+x)$.

A third natural distance is the *Earth Mover distance* (EM), where $\delta_{EM}(\mathbf{s}', \mathbf{s}) \leq x$ if \mathbf{s}' can be attained from \mathbf{s} by changing the vote of at most x voters (similar to the ℓ_1 norm, but with the constraint that the number of votes remains the same).

We can also consider more elaborate ways of defining possible world states, e.g., where a voter is optimistic or pessimistic regarding the score of his favorite candidates; EM distance where voters only transfer their votes to higher-ranked candidates,⁵ and so on. For ease of presentation, we will consider the ℓ_1 norm, unless explicitly stated otherwise.

⁵A similar assumption, together with a variation of the multiplicative metric, was used in (Reyhani et al., 2012) to determine the set of possible winners.

4.4. Strategic voting and equilibria

Let $g_i : M^n \rightarrow M$ be a *response function*, i.e. a mapping from voting profiles to actions (which implicitly depends on the preferences of voter i). Any set of response function $(g_i)_{i \in N}$ induces a (deterministic) dynamic in the normal form game corresponding to a particular preference profile under Plurality. In particular, it determines all equilibria of this game, which are simply the states \mathbf{a} where no voter has a response that changes the state. We refer to the response function of a voter as her *type*. We emphasize that the set of voting equilibria depends only on voters' preferences and response functions, and not on whether they vote iteratively or simultaneously.

Definition 3. Let N be a set of voters with response functions $(g_i)_{i \in N}$. A *voting equilibrium* is a state \mathbf{a} , where $a_i = g_i(\mathbf{a})$ for all $i \in N$.

We next define the primary response function that strategic voters in our model apply, which is based on local dominance.

Definition 4. A *strategic voter of type r* (or, in short, an r voter) acts as follows in state \mathbf{a} . Let $D \subseteq M$ be the set of candidates that $S_i(\mathbf{a}, r)$ -dominate a_i . If D is non-empty, then i votes for his most preferred candidate in D . Formally, $g_i(\mathbf{a}) = \operatorname{argmin}_{d \in D} Q_i(d)$ if $D \neq \emptyset$, and $g_i(\mathbf{a}) = a_i$ otherwise.

We refer to r (or r_i if types differ) as the *uncertainty parameter* of the voter. We denote such a strategic step by $a_i \xrightarrow{i} a'_i$, where $a'_i = g_i(\mathbf{a})$. We observe that:

- If a'_i $S_i(\mathbf{a}, r)$ -dominates a_i , then a'_i $S_i(\mathbf{a}, r')$ -beats a_i , for any $r' \geq r$.
- For $r = 0$, the voter knows the current voting profile exactly, and thus his response function is simple best-response, as in (Meir et al., 2010).
- For $r = n$, a voter does not know anything about the current voting profile. Thus an action a'_i locally dominates a_i if and only if it weakly *globally dominates* a_i . Note that this typically means a_i is i 's least preferred candidate.⁶

Different definitions of strategic responses (distance metrics, value of r) may induce different sets of voting equilibria. However, the assumption that i votes for the most-preferred candidate in D is irrelevant to the set of induced equilibria. The following is an immediate observation.

⁶This may not hold if $n < m$. E.g. if i is the only voter, then her top preference globally dominates all other candidates.

Proposition 1. *Let N be a set of voters with preferences \mathbf{Q} and following Def. 4 (voters may be of heterogeneous types). A voting profile \mathbf{a} is a voting equilibrium, if and only if no voter votes for a locally dominated candidate. Formally, if $\forall i \in N, \nexists a'_i \in M$, such that a'_i $S_i(\mathbf{a}, r_i)$ -dominates a_i .*

We assume of course, that the same parameters are used for defining the dominance relation and the strategic response of each voter. For example, under Definition 4 with $r = 0$, a strategic move coincides with best-response, and voting equilibria coincide with pure Nash equilibria of the base game \mathbf{Q} .

5. Convergence with Strategic Voters

For any $w \in \mathbb{N}$, let $H_w(\mathbf{s}) \subseteq M$ be the set of candidates that need exactly w more votes to become the winner. Thus $H_0(\mathbf{s}) = \{f(\mathbf{s})\}$, $H_1(\mathbf{s}) = \{c : s(f(\mathbf{s})) > s(c) \geq s(f(\mathbf{s})) - 1\}$ (either have the same score as the winner and lose by tie-breaking, or c wins in the tie-breaking but have one vote less), etc. Let $\overline{H}_w(\mathbf{s}) = \bigcup_{w' \leq w} H_{w'}(\mathbf{s}) = \{c : s(c) \geq s(f(\mathbf{s})) - w\}$.

5.1. Strategic responses and possible winners

We say that candidate c is a *possible winner for i in state \mathbf{a}* if there is an accessible state where c wins. Formally, $W_i(\mathbf{a}, r) = \{c \in M : \exists \mathbf{s}' \in S_i(\mathbf{a}, r), f(\mathbf{s}', c) = c\}$.

In contrast with $\overline{H}_w(\mathbf{s})$, the definition of $W_i(\mathbf{a}, r)$ depends on the identity of the voter, and not only on her type.

We first show that in every strategic response, a voter always votes for her favorite possible winner.

Lemma 2. *Consider a strategic response $a_i \xrightarrow{i} a'_i$ s.t. $a_i \notin W_i(\mathbf{a}, r)$. Then $a'_i = \operatorname{argmin}_{c \in W_i(\mathbf{a}, r)} Q_i(c)$.*

The lemma holds for any ℓ_d norm, $d \geq 1$, the multiplicative distance, and EM distance.

Proof. Consider the set of candidates D in Def. 4. Since $a'_i \in D$, it is non-empty. We will show that (I) $D \subseteq W_i(\mathbf{a}, r)$; and (II) $b = \operatorname{argmin}_{c \in W_i(\mathbf{a}, r)} Q_i(c)$ is in D . This implies $b = a'_i$.

For (I), every candidate in D in particular $S_i(\mathbf{a}, r)$ -beats a_i , and therefore must be a possible winner (transferring i 's vote from one non-possible winner to another does not change the outcome in any accessible state).

For (II), first note that since $a'_i \in D$, there is a state $\mathbf{s}' \in S_i(\mathbf{a}, r)$ where i prefers $f(\mathbf{s}', a'_i) = a'_i$ over $f(\mathbf{s}', a_i) = a_i$, thus $a'_i \succ_i a_i$. Also, by (I) $a'_i \in$

$W_i(\mathbf{a}, r)$ and thus $b \succeq_i a'_i \succ_i a_i$. Thus a_i cannot $S_i(\mathbf{a}, r)$ -beat b . It remains to prove that b $S_i(\mathbf{a}, r)$ -beats a_i (and thus locally dominates it).

We consider two cases. Suppose first that for all $\mathbf{s}' \in S_i(\mathbf{a}, r)$, $f(\mathbf{s}', a_i) = b$. Then for any $c \in M$, and any \mathbf{s}' , $f(\mathbf{s}', c) \in \{b, c\}$. Moreover, by definition of b , if $f(\mathbf{s}', c) = c$ then $c \preceq_i b$. Thus no candidate $S_i(\mathbf{a}, r)$ -beats a_i , and $D = \emptyset$. Thus suppose there is a state $\mathbf{s}' \in S_i(\mathbf{a}, r)$, $f(\mathbf{s}', a_i) = c \prec_i b$, and there is also a state (since b is a possible winner) $\mathbf{s}'' \in S_i(\mathbf{a}, r)$, $f(\mathbf{s}'', b) = b$. It can be verified for each of the metrics we consider (ℓ_d norms, multiplicative, EM) that there must be a state $\mathbf{s}^* \in S_i(\mathbf{a}, r)$ where $f(\mathbf{s}^*, b) = b$ but $f(\mathbf{s}^*, a_i) = c$. For the ℓ_1 norm, we get \mathbf{s}^* from \mathbf{s}' by adding votes to b until c, b are tied (or until there is a difference of one for c , if b beats c in \overline{Q}).⁷ Thus b $S_i(\mathbf{a}, r)$ -beats a_i . \square

Lemma 2 does not mean that our dynamics coincides with “always vote for the most preferred possible winner”. It only holds when the current vote of i is *not* a possible winner, and when there are at least two possible outcomes. If there is only one possible outcome, i will not move (which makes sense). Also, if a_i is a possible winner, then typically i will not move, and if he does move it may be to a non-possible winner. For example, if a_i is the least-preferred possible winner, than *any* other candidate locally dominates a_i . We will later see that while these situations are hard to analyze, they can be avoided in some paths to equilibrium.

Threshold for possible winners. We continue with the following lemma, which shows that for some simple metrics, the set of possible winners is exactly all candidates whose score is above a certain threshold. Denote by $f_i^* = f(\mathbf{a}_{-i})$ the candidate that would win if i would not vote.

Lemma 3. *Each of the metrics δ from $(\ell_1, \ell_\infty, \text{multiplicative})$ induces a function $\beta = \beta_{\delta, r} : \mathbb{N} \rightarrow \mathbb{N}$, where*

- *For every $\mathbf{a}, i \in N$, if $a_i \neq f_i^*$, then $W_i(\mathbf{a}, r) = \{c : s_{\mathbf{a}_{-i}}(c) \geq \overline{Q} \beta(s(f_i^*))\}$. That is, possible candidates are all those whose score is above the threshold, which is a function of the score of the winner.⁸*
- *$\beta(s)$ is weakly increasing in s .*

⁷For other ℓ_d norms and multiplicative distance, we can consider any path of states between $\mathbf{s}', \mathbf{s}''$ that is contained in $S_i(\mathbf{a}, r)$, by gradually removing votes from c and from other candidates, and adding votes to b . The critical state \mathbf{s}^* will be along this path. For EM distance some paths may fail if we transfer votes directly from c to b , but we can use a path where we transfer votes first from winners to a_i , and then from a_i to b .

⁸We break ties with $\beta(s(f_i^*))$ as if we break ties with f_i^* .

In particular, for δ_{ℓ_1} , $W_i(\mathbf{a}, r) = \overline{H}_{r+1}(\mathbf{s}_a)$ for any i s.t. $a_i \notin \overline{H}_{r+1}(\mathbf{s}_a)$.⁹

A similar result can be proved for EM and other ℓ_d norms, but the threshold would depend on the score of all candidates and not just the winner.

Proof for the ℓ_1 metric. Consider first the ℓ_1 norm. We set $\beta(s) = \beta_{\ell_1}(s) = s - r - 1$. Clearly, f_i^* is a possible winner and $s(f_i^*) > \beta(\mathbf{a})$. Assume $c \neq f_i^*$, then $s(c) \not\geq \beta(s(f_i^*))$. If $s(c) > \beta(s(f_i^*))$, consider the state $\mathbf{s}' \in S_i(\mathbf{a}, r)$ where c has r additional votes. We have that $s'(c) + 1 = s(c) + r + 1 > \beta(s(f_i^*)) + r + 1 = s(f_i^*)$, thus $f(\mathbf{s}', c) = c$. In contrast, if $s(c) < \beta(s(f_i^*))$, then in any $\mathbf{s}' \in S_i(\mathbf{a}, r)$, $s'(f_i^*) - s'(c) > 1$ and thus c cannot win. Finally, since $a_i \notin W_i(\mathbf{a}, r)$, then $s_{\mathbf{a}}(c) = s_{\mathbf{a}_{-i}}(c)$ for all $c \in W_i(\mathbf{a}, r)$. Thus $W_i(\mathbf{a}, r) = \{c : s(c) \geq \beta_{\ell_1}(s(f_i^*))\} = \{c : s(c) \geq s(f_i^*) - (r + 1)\} = \overline{H}_{r+1}(\mathbf{s})$. \square

Note that Lemmas 3 and 2 together entail that the strategic decision of the voter is greatly simplified from both a behavioral and a computational perspective. There is no need to consider all possible world states. Only to check which candidates are sufficiently close to the winner in terms of their prospective score, and select the one that is most preferred.

5.2. Existence of equilibrium and convergence from the truthful state

In what follows, we will only consider the ℓ_1 norm for simplicity. However most results hold for other metrics as well. We also highlight that as in (Meir et al., 2010), we allow any number of non-strategic voters as part of the input. To allow continuous reading, some of the proofs were deferred the appendix.

Best-response graphs and schedulers. Given a game, any dynamic induces a directed graph whose vertices are the states of the game (M^n in the case of the Plurality game). There is an edge from a state \mathbf{a} to a state \mathbf{a}' , if (1) \mathbf{a}, \mathbf{a}' differ only by the action of one player i ; and (2) $g_i(\mathbf{a}) = a'_i$. We call this graph the *best-response* graph. We can similarly create a *group best-response graph* where edges correspond to the (non-coordinated) actions of subsets of agents. That is, there is an edge $(\mathbf{a}, \mathbf{a}')$ iff there is a subset $N' \subseteq N$ s.t. $a'_i = g_i(\mathbf{a})$ for all $i \in N'$, and $a'_i = a_i$ for $i \notin N'$.

A *scheduler* selects which voters play at any step of the game. A scheduler ϕ is a *singleton scheduler* if it always selects a single voter, and otherwise it is a *group scheduler*.¹⁰ We assume that the order of players is determined by

⁹Similarly, for δ_{ℓ_∞} , $W_i(\mathbf{a}, r) = \overline{H}_{2r+1}(\mathbf{s}_a)$ for any i s.t. $a_i \notin \overline{H}_{2r+1}(\mathbf{s}_a)$. See appendix for the full proof.

¹⁰We emphasize that a group of voters moving at the same time *does not* coincide with a coalitional manipulation. Each player acts as if he is the only one moving, and it may well be that the result of a move by a group is worse for all of its members.

an arbitrary singleton scheduler (see (Apt and Simon, 2012)). Equivalently, a scheduler can be thought of as a *tie-breaking* mechanism for turns, when more than one voter or set of voters have a strategic move.

We next show that when starting from the truthful state, a singleton scheduler guarantees convergence to an equilibrium. In particular, an equilibrium must exist. Proposition 4 also follows from more general convergence results that we will show later. However, we provide a simple and detailed proof that reveals the natural structure of the equilibrium that is reached. In the appendix, we show how to extend the proof to other distance metrics (Theorem 10).

Theorem 4. *Suppose that all voters are of type r . Then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, for any singleton scheduler, they will converge to an equilibrium in at most $n(m - 1)$ steps.*

Proof. If the truthful state \mathbf{q} is stable, then we are done. Thus assume it is not. Let \mathbf{a}^t (and \mathbf{s}^t) be the voting profile after t steps from the initial truthful vote $\mathbf{a}^0 = \mathbf{q}$. Let $a_i \xrightarrow{i} a'_i$ be a move of voter i at state $\mathbf{s} = \mathbf{s}_{\mathbf{a}^t}$ to state $\mathbf{s}' = \mathbf{s}_{\mathbf{a}^{t+1}}$. We claim that the following hold throughout the game. Recall that by Lemma 3, \overline{H}_{r+1} is the set of possible winners for all voters who are voting for other candidates.

1. $a_i \notin \overline{H}_{r+1}(\mathbf{s}')$, i.e., once a candidate is deserted, it is no longer a possible winner.
2. $a'_i \prec_i a_i$, i.e., voters always “compromise” by voting for a less-preferred candidate.
3. $\max_{a \in M} s'(a) \geq \max_{a \in M} s(a)$, i.e., the score of the winner never decreases.
4. $\overline{H}_{r+1}(\mathbf{s}') \subseteq \overline{H}_{r+1}(\mathbf{s})$, i.e., the set of possible winners can only shrink.

We prove this by a complete induction.

1. If this is the first move of i then $a_i = q_i$. Otherwise, by Lemma 2, a_i is the most-preferred candidate of i in $\overline{H}_{r+1}(\mathbf{s}^{t'})$ where t' is the time when i last moved. By induction on (4), $\overline{H}_{r+1}(\mathbf{s}) \subseteq \overline{H}_{r+1}(\mathbf{s}^{t'})$. So if $a_i \in \overline{H}_{r+1}(\mathbf{s})$, it must be the most-preferred candidate in the set. Assume, towards a contradiction, that $a_i \in \overline{H}_{r+1}(\mathbf{s})$; then there is a state $\hat{\mathbf{s}} \in S_i(\mathbf{a}, r)$ where i is pivotal between a_i and $f(\mathbf{a}) \prec_i a_i$. For any $c \neq a_i$, $f(\hat{\mathbf{s}}, c) = f(\mathbf{a})$, and in particular for $c = a'_i$. Therefore a_i $S_i(\mathbf{a}, r)$ -beats a'_i , which means a'_i does not $S_i(\mathbf{a}, r)$ -dominate a_i . A contradiction.




Figure 1: An example of 18 voters voting over candidates $\{a, b, c, d, e\}$. The top left figure shows the initial (truthful) state of the game. The letter inside a voter is his second preference. The dashed line marks the threshold β_{l_1} of possible winners for voters of type $r = 2$. Thus candidates on or above the threshold are the set $\overline{H}_3(\mathbf{a})$. Candidates that are on the dashed line (in $H_3(\mathbf{a})$) are considered possible winners only by voters that do not currently vote for them. Note that due to tie breaking it is not the same for all candidates. For example, since a beats b in tie-breaking, b needs 2 more votes to win in the initial state. In the next two figures we can see voters leaving their candidates (who are not possible winners for them) to join one of the leaders. The last figure shows an equilibrium that was reached. Note that although both of a, c are possible winners for the last supporter of b , he has no strategic move. A different equilibrium may have been reached with a different scheduler.

2. If this is the first move of i then this is immediate. Otherwise, by induction on Lemma 2 and (4), if $a'_i \succ_i a_i$, then i would prefer to vote for a'_i in his previous move, rather than for a_i .
3. As in (1), if i votes for $a_i = f(\mathbf{a}) \in \overline{H}_{r+1}(\mathbf{a})$, then a_i is i 's most-preferred possible winner. Thus it cannot be locally dominated by any other candidate.
4. Since by (3) the score of the winner never decreases, the only way to expand \overline{H}_{r+1} is to add a vote to a candidate not in \overline{H}_{r+1} . By Lemma 2 this never occurs.

Finally, by property (2), each voter moves at most $m - 1$ times before the game converges. \square

The proof not only shows that an equilibrium exists, it also describes exactly the way in which such equilibria are reached from the truthful state. There is always a set of “leaders” (\overline{H}_{r+1} in the case of the ℓ_1 norm). Strategic voters vote for their favorite candidate in this set, if their current candidate is not a possible winner. At some point candidates may “drop out” of the race as their gap from the winner increases, and the set \overline{H}_{r+1} shrinks. This continues; in the reached equilibrium, all strategic voters vote for their best possible winners (which is in \overline{H}_{r+1}). As we will see next, the only case where there are voters $N' \subseteq N$ not voting for a possible winner, is when the gap between the winner and the runner-up is exactly $r+1$, and all of N' prefer the current winner. In particular, the gap can never grow larger than $r+1$, and thus there are always at least two leaders whenever there is some strategic interaction (see Figure 1).

Lemma 5. *Under the conditions of Theorem 4, either \mathbf{q} is stable, or in every state \mathbf{s}^t we have $|\overline{H}_{r+1}(\mathbf{s}^t)| > 1$. Also, in the last state \mathbf{a} either $|\overline{H}_r(\mathbf{s}_{\mathbf{a}})| = 1$, or all voters vote for possible winners. Any voter voting for $c \notin \overline{H}_{r+1}(\mathbf{s}^t)$ prefers $f(\mathbf{a})$ over any other candidate in $\overline{H}_{r+1}(\mathbf{s}^t)$.*

5.3. Convergence under broader conditions

To show robust convergence results, there seem to be two main extensions. First, we would like the game to converge from any initial state, and not just from the truthful one. Second, we would like convergence to occur even if more than one voter moves between states. In other words, we would like to see convergence under any group scheduler.

Unfortunately, under arbitrary group schedulers, convergence is not guaranteed even from the truthful state. A simple example for $r = 0$ appears in (Meir et al., 2010), where there are two runner-ups that are preferred to the

winner by all of their supporters, but the supporters fail to coordinate on a runner-up to promote. We conjecture that as in the case of best-responses ($r = 0$), a *singleton scheduler* would guarantee convergence from any initial state.

We show that if we make two mild restrictions on the scheduler, then convergence from any state is guaranteed even for groups. We say that a step $a \xrightarrow{i} a'$ is of type 1 if $a' \prec_i a$, and type 2 if $a' \succ_i a$. We call type 1 steps *compromise steps*, and type 2 *opportunity steps*.

Proposition 6. *Suppose that all voters are of type r . Consider any group scheduler such that (1) any voter has some chance of playing as a singleton (i.e., this will occur eventually); (2) the scheduler always selects (an arbitrary subset of) voters with type 2 moves, if such exist. Then convergence is guaranteed from any initial state after at most $O(nm)$ singleton steps occur.*

The proof shows that in particular, for singleton schedulers there is a path of best-responses from any state to an equilibrium (any singleton path where type 2 steps precede type 1 steps). Starting from the truthful state is a special case, where there are no type 2 moves before the first type 1 move. The assumption that type 2 moves are played first can be justified to some extent, since type 1 moves are “compromises” and thus voters may be more reluctant to carry them out.

5.4. Truth-bias and Lazy-bias

The basic strategic behavior from Definition 4 still allows for some counter intuitive actions. For example if there is only one possible winner, then a voter will not change her vote regardless of how much she dislikes her current action. In particular we can still construct an equilibrium where all voters vote for their least preferred candidate!

However, the notion of local dominance is very flexible, and allows us to define more subtle and plausible behaviors. Specifically, by adding a negligible utility ϵ to a favorite action, such as truth-telling or abstaining, we get that this action locally dominates any other action where the outcome is the same (that is, the same in all accessible world states). We can thus define *truth-biased* or *lazy* voters, who prefer to tell the truth or to abstain whenever they do not see themselves as pivotal. We highlight that the local neighborhood considered by the voter when deciding whether to cast a strategic vote and when applying truth-bias/lazyness, is not necessarily the same neighborhood.

Definition 5. A strategic truth-biased voter of type (r, k) (or, in short, a $T(r, k)$ voter) acts as follows in state \mathbf{a} .

1. (strategic move) Let $D \subseteq M$ be the set of candidates that $S_i(\mathbf{a}, r)$ -dominate a_i . If D is non-empty, then i votes for his most preferred candidate in D .
2. if a_i $S_i(\mathbf{a}, k)$ -beats q_i , then i keeps current vote a_i .
3. (truth bias move) otherwise, i moves to q_i .

A *strategic lazy voter of type (r, k)* (an $L(r, k)$ voter) can be similarly defined, replacing q_i with the action \perp (abstain).

We make the following observations.

- A $T(r, n)$ or $L(r, n)$ voter is just an r voter. This is since truth/lazy-bias will never be applied.
- $T(0, 0)$ voters follow the truth-bias model in (Thompson et al., 2013; Obraztsova et al., 2013).
- $L(0, 0)$ voters follow the (simultaneous) “lazy” model of Desmedt and Elkind (2010).

Intuitively, the parameters r and k reflect two uncertainty thresholds that determine how much a voter is inclined to make a strategic response, and how much to keep his current vote. We argue that it is natural to assume $k > r$, which entails that a voter requires a lower uncertainty level in order to make a new strategic step, than to merely keep his current strategic vote. From a behavioral perspective, such an assumption accounts for *default-bias*: decision makers have a higher tendency to stay with their current decision, than to adopt a new one (Kahneman et al., 1991). Lower values of k (i.e., closer to r) correspond to a stronger truth-bias, whereas higher values correspond to a stronger default bias.

Proposition 7. *Suppose that each voter i is of type $L(r, k_i)$ or $T(r, k_i)$, where $k_i > r$. Then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, they will always converge to an equilibrium in at most $3nm$ steps.*

5.5. Other considerations

General voting equilibria. What is the structure of voting equilibria that are not obtained by starting from the truthful state? They can be more diverse, but no voter votes for a locally dominated candidate (Prop. 1).

Corollary 8. *In any voting equilibrium \mathbf{a} , for any set of heterogeneous voters (i.e. voter i is of type $T(r_i, k_i)$ or $L(r_i, k_i)$), it holds that:*

1. *Every voter is either truthful, or votes for a candidate in $\overline{H}_{k_i}(\mathbf{a})$.*
2. *No voter votes for his least-preferred candidate in $W_i(\mathbf{a}, r_i)$.*

In particular, when k_i is not too high, this rules out “crazy” equilibria where voters vote for their least preferred candidate. It is possible, though, that all voters vote for their second-least-preferred possible winner, or their second-least-preferred candidate (consider a profile where candidates a, b are ranked last by all voters, where roughly half rank a above b).

Unique best-response. Meir et al. (2010) argued for a “unique best-response”. They enforced a requirement that voters only move to candidates that, as a result, become winners. We show that when $r > 0$, the notion of local dominance provides some natural justification for this requirement.¹¹

Observation 9. *Let $r > 0$, $S = S_i(\mathbf{a}, r)$. Suppose that c is the most preferred possible winner for i , and b is not a possible winner for i . Then c $S_i(\mathbf{a}, r)$ -dominates b .*

Intuitively, while voting for b may improve the outcome for i (over a_i) in some world states, voting for c will be at least as good in those states, and strictly better in the states where i is pivotal between c and another possible winner.

6. Simulations of strategic voting

We explore via extensive simulations how employing local-dominance affects the result of the voting process. These simulations have two primary goals. First, we want to understand better the effect of different parameters on the technical level (for example, how long does it take to reach an equilibrium if we vary the uncertainty level?). More importantly, we use simulations to test the properties of our strategic model with respect to the desiderata listed on Section 2. For example, what is its discriminative power, is it robust to small changes, and whether it replicates common phenomena. See Section 7.1 for a summary of our findings in light of the desiderata.

Preferences. As the outcome depends not just on our strategic model but also, and perhaps mainly, on the preference profile, we need to specify appropriate distributions on preferences for our simulations. We generate preference profiles from a set of distributions which have been examined in the

¹¹We thank Greg Stoddard for a discussion leading to this observation.

research literature, with a focus on distributions that are claimed to resemble preferences of human societies: The Uniform (or impartial culture) distribution; a uniform Single-peaked distribution; Polya-Eggenberger Urn model (with 2 urns and with 3 urns); a Riffle distribution; and Plackett-Luce distribution. Urn models were particularly designed to resemble preference structures in human societies, whereas in Plackett-Luce distributions each voter is assumed to have a noisy signal of some ground truth. See Appendix B.1 for details.

Methods. We generated profiles from all distribution types for various numbers of voters and candidates, which resulted in 108 distinct distributions. From each distribution we sampled 200 instances.¹² Then, we simulated strategic voting on each instance varying the distance metric (ℓ_1 , multiplicative), the voters’ types (basic, truth-biased, lazy) and the uncertainty parameters r and k .

We simulated voting in an iterative setting, where voters start from an initial state, and then iteratively make strategic moves until convergence. Simulations all started from the truthful outcome, except for one batch on which we will elaborate later. We repeated each simulation 100 times (as the scheduler may pick a different path each time), and collected multiple statistics on the equilibrium outcomes. See Appendix B.2 for details. All of the collected data can be downloaded from <http://tinyurl.com/k2b775e>.

6.1. Results

Meaningful parameters and robustness. We observed that results with the multiplicative metric were generally very similar to those with the additive (ℓ_1) metric. Also, while voting with lazy and truth-biased voters resulted in a somewhat different dispersion of the votes in equilibrium, all of the properties that we measured remained largely the same, regardless of the value of k . We conclude that these parameters have little effect on the model when the initial state is truthful, and focus on results for voters with additive distance metric and without truth- or lazy-bias.¹³

The choice of scheduler type also turned out to be an immaterial one. While a group scheduler typically converged much faster, there was only

¹²We also used three datasets from German pre-election polls, with 100 voters, 3 candidates, and no sampling, as well as all 225 currently available full preferences from PrefLib (<http://preflib.org>).

¹³With lazy-bias simulations typically took more steps to converge.

a negligible difference in the equilibrium outcomes for the vast majority of preference profiles.

The most meaningful parameter in the simulations was the uncertainty level r . As we vary the value of r from 0 to 15, there is an increase and then a decrease in the amount of strategic behavior, with a “peak value” for r . We can see the effect of more strategic behavior by looking at the number of steps to convergence (Figure 2), the higher dispersion of equilibrium states, and the (lower) agreement with the Plurality winner (Figure B.8 in appendix). This pattern makes sense, as with low r the voter knows the current state exactly, and often realizes he is not pivotal. As uncertainty grows the voter considers himself pivotal more often, but beyond the peak r uncertainty is sufficiently large for all voters to believe that their truthful vote is also a possible winner (and then the initial state is stable).

This pattern repeats in all 108 distributions. The effect of r and in particular its peak value are determined mainly by the type of the distribution and the number of voters, where the peak r increases with n (for distributions with $n = 10$, r typically peaks at 0). The number of candidates may affect the strength of the strategic effect, but not the peak r .




Figure 2: Average number of steps to reach equilibrium for the 2-urn distribution ($NumStep$), as a function of the uncertainty parameter r . Each group of lines marks a different number of voters, and peaks at a different value of r . Recall that for $r = n$ (denoted by ‘max’), we have no strategic steps. This serves as our baseline.

Quality of winner. In the Plackett-Luce distribution, the quality of a winner can be determined according to its rank in the ground truth used to generate

the profile. In the other distributions there is no notion of ground truth, and hence we measured how often the equilibrium winner agreed with the (truthful) winner of another common voting system, which takes the entire preference profile into account (Borda, Copland, Maximin). We also measured how often the winner was a Condorcet winner (out of cases where one exists), and the social welfare of voters (assuming Borda utilities).

According to Borda, Copland, Condorcet consistency, social welfare (see Figure 3) and the ground truth, a clear pattern was observed almost invariably across all distributions. As strategic activity increases, so does the winner quality.¹⁴ Best winner quality is attained at peak r or very close to it (see Figures B.9, B.11 in the appendix).

In particular, these results are interesting for the Single-Peaked profiles. In such profiles there is always a Condorcet winner, which is the median candidate. As voters strategize more under Plurality, they in fact get closer to the outcome of the strategy-proof median mechanism.




Figure 3: The increase in voters’ social welfare, compared to the truthful Plurality outcome.

Duverger Law. The positive effect of more strategic behavior on the concentration of votes was remarkably clear across all distributions. As r gets closer to the peak value, over 75% of the voters (all voters in some distributions) end up voting for only two candidates. This holds both in distributions like Uniform and Plackett-Luce where no candidate initially has a strong advan-

¹⁴There was typically no higher agreement with the Maximin winner. Also, in the Urn models strategic behavior did not in general improve consistency with Borda.

tage (see Figure B.12 in appendix), and in distributions like 2-Urn where there are two leading candidates to begin with (Figure 4).




Figure 4: The fraction of simulations in 2-Urn distribution in which *all* voters ended up voting for only two candidates (*TotalDuverger*).

Real preference datasets. In general, all of our empirical findings were replicated on the real preference data, but since each instance has a different number of voters and candidates, results are more qualitative.

In all three German election datasets we observed similar patterns as above: nearly all supporters of the third candidate deserted it to join one of the leaders. In two datasets this did not change the identity of the Plurality winner. In the third dataset the strategic behavior (for any r between 1 and 30) replaced the Plurality winner with the Condorcet winner.

Similarly, in most of the PrefLib datasets there was a clear winner, and thus there were none-to-few strategic moves. Votes were typically already quite concentrated for the two leaders in the initial truthful profile, but this concentration increased with strategic activity (Duverger’s law). In the few instances where the identity of the winner changed, it usually replaced the Plurality winner with the Condorcet winner.

Non-truthful starting profile and Discriminative power. We ran a batch of simulations starting from a voting profile chosen uniformly at random. While we observed voters sometimes switching back and forth between candidates, and despite having no formal guarantee of convergence, all simulations eventually converged to an equilibrium.

We generally observe that the dependency of various attributes in r is more complex (there is no clear “peak r ”), but r and n are still the most significant factors. For example, we see much more strategic activity with $r = 0$ as in a random state there are more likely to be many pivotal voters than in the truthful state.

Other properties observed above such as Duverger’s law, and an increase in winner quality are replicated when starting from a random profile (see e.g. Figure B.10 in appendix).

More importantly, simulations with random initial states enable us to test the discriminative power of the model.¹⁵ Without strategic behavior, we would not expect any candidate to win in a large fraction of the instances (e.g., in the uniform distribution every candidate should win in $\sim 1/m$ of the instances). However when voters are strategic we get that most of the simulations select the same winner regardless of the initial state, which indicates high discriminative power (Figure 5).




Figure 5: The fraction of simulations (out of 100 random initial states for each preference profile) that ended with the same winner (*WinnerConsistency*). Note that around peak r we have 75% – 85% of the starting points leading to the same winner, regardless of the number of candidates. Very similar patterns were observed for the Urn models and Plackett-Luce.

Diverse population. Finally, we repeated some of our simulations with heterogeneous voters, where r_i are sampled uniformly i.i.d. from $\{0, 1, \dots, n/m\}$. Despite the fact that our convergence proofs do not cover heterogeneous

¹⁵Looking only on simulations from the truthful profile does not mean much in that respect. Note for example that with $r = n$, we would get “perfect” discriminative power as the truthful Plurality winner is always selected.

populations, convergence was just as robust. Not only that all simulations converged, typically all invariants that we proved for the homogenous case (e.g., that a voter always compromises for less desirable candidates) also hold in the diverse case.¹⁶

The diverse simulations replicated nearly all the patterns of strategic voting across all distributions. Notably, although we used the same simple distribution of r_i values in all simulations, effects of strategic behavior were always approximately as strong as in the peak r value of every profile distribution and across most measured properties.

Winner quality was also generally comparable to peak r .¹⁷

Regarding Duverger’s law, while the number of votes to the top 2 candidates was generally quite similar to the one in peak r (sometimes even higher), with diverse population there were much *fewer* instances where *only* two candidates received votes. See Figure 6. Looking at a typical equilibrium profile reveals that it has a much more “natural” dispersion, with many voters voting for the two leaders, but also some voters (with either very high or very low uncertainty values) voting for other candidates.

7. Discussion

In (Abramson et al., 1992), sophisticated (strategic) voting based on expected utility maximization is defended on the grounds that it “...is a simplification of reality that seeks to capture the most salient features of actual situations. Many voters may see some candidates as having real chances of winning and others as likely losers, and they may weigh these perceptions against the relative attractiveness of the candidates.”

Our theory is also a simplification of reality, and applies similar logic to explain and justify strategic voting. However, the local-dominance approach allows voters to take into account both “chances of winning” and “relative attractiveness”, without regressing to probabilistic calculations and expected utility maximization.

7.1. The model and the desiderata

We summarize by showing how model of local dominance answers to the desiderata we presented in Section 2.

¹⁶There were as few as 4 violations of the invariant out of $\sim 120,000$ strategic steps.

¹⁷Interestingly, the Plackett-Luce distribution is an exception, where diverse population led to degradation in the winner quality according to the ground truth, but not according to the other measures.




Figure 6: In the top figure we can see that with diverse population, votes were just as concentrated as with fixed population with peak r , across all distributions. The bottom figure shows that with fixed r , this concentration is due to many instances where only two candidates get votes, while this is not the case with diverse population.

- Looking at *theoretical criteria*, our model is grounded in traditional game-theoretic concepts: voters are trying to maximize their utility, and results are in equilibrium. Further links to decision theory and classical notions of rationality are detailed in Section 7.2.

When all voters are of the same type, an equilibrium always exists, and convergence of local-dominance dynamics is guaranteed under rather weak conditions. Our simulations show existence and convergence even without these conditions, and demonstrate high discriminative power. The model is broad enough to encompass different scenarios such as simultaneous, sequential and iterative voting, and to account for behaviors such as truth-bias and lazy-bias. Our definitions could be easily extended to other positional scoring rules, although it is an open question whether our results would still hold.¹⁸ Furthermore, as Proposition 8 shows, preposterous equilibria are unlikely.

¹⁸Extensions to other voting rules are simple with metrics like EM distance, but may be ill-defined with other metrics.

- As argued above, voters in our model fit the *behavioral criteria* we posed, as they avoid complex complex computations. Moreover, as Lemma 3 shows, voters do not even need to consider the entire space of possible states, but merely to check which candidates have sufficient score to become possible winners. Our informational assumptions are rather weak and plausible, as we argue in the end of Section 4.1.
- Regarding the *scientific criteria*, once we set the distance metric, every voter can be described by a single parameter (two in the case of lazy or truth-biased voters), which has a clear interpretation as her certainty level. Our extensive simulations demonstrate robustness to the order in which voters play (including whether they act simultaneously or not), and that changing the parameters results in a rather smooth transition. Simulations also show that the model replicates patterns that are common in the real world such as the Duverger Law, and resulting equilibria, especially with diverse population, seem reasonable. Experimental validation as outside the scope of this work.

Finally, it is shown that strategic behavior yields a better winner for the society according to various measures of quality (compared to the truthful Plurality winner).

7.2. Epistemic foundations and rationality

We can phrase dominance relations in terms of modal logic. Consider a Kripke structure over states where S are the states accessible from $\mathbf{s} = \mathbf{s}_{\mathbf{a}}$. Then “ b_i S -beats a_i in state \mathbf{s} ” can be written as $\mathbf{s} \models \diamond(f(b_i, \mathbf{s}) \succ_i f(a_i, \mathbf{s}))$. Similarly, “ b_i S -dominates a_i ” means $\mathbf{s} \models \Box(f(b_i, \mathbf{s}) \succeq_i f(a_i, \mathbf{s})) \wedge \diamond(f(b_i, \mathbf{s}) \succ_i f(a_i, \mathbf{s}))$ (b_i is necessarily at least as good, and possible a better action than a_i). We note that $S_i(\mathbf{a}, k)$ defines a Kripke structure that is reflexive and symmetric but non-transitive.

A common semantic interpretation of the modal operator $\Box P$ is “ P is known”.¹⁹ According to this, we can naturally interpret “ b_i S -beats a_i in state \mathbf{s} ” as “in state s , i does not know that voting for a_i is at least as good as voting for b_i ”. Thus “ b_i S -dominates a_i in state \mathbf{s} ” means that i knows that b_i is at least as good as a_i , but does not know that a_i is at least as good as b_i .

¹⁹An alternative notation $K_i P$ is sometimes used for the statement “ P is known to agent i ”. See, for example, (Aumann, 1999).

In our model, $S(\mathbf{a})$ is not uniquely defined, and in fact even the same voter uses both $S_i(\mathbf{a}, r)$, $S_i(\mathbf{a}, k)$, where $r < k$. Since $S_i(\mathbf{a}, r) \subseteq S_i(\mathbf{a}, k)$, a straightforward extension of the epistemic interpretation is to add *certainty levels*, where a larger (in terms of containment) set of accessible states indicates lower certainty.

Local dominance and rationality. According to the standard non-Bayesian incomplete information model (due to Aumann (1995; 1999)), a player i playing strategy a_i at some world state π is *rational*, if there is no other strategy a'_i that yields a same or better outcome in all states accessible from π (and in some states strictly better).

In other words, rationality under strict uncertainty according to Aumann simply means that players avoid locally dominated strategies. Voting equilibria in our model are therefore rational (Prop. 1). Our model is more specific in that it specifies a particular dynamic of how voters act when their current strategy is dominated.

Another difference is that in Aumann’s models the accessibility relation is a partition, and in particular transitive. Other papers such as (Bicchieri and Antonelli, 1995) do not make any assumptions on the accessibility relation other than consistency. In our model the relation is based on distance, and in particular it is non-transitive (if π is close to π' , and π' is close to π'' , then it may not hold that π is close to π'').

Local dominance and voting. Dominance within a restricted set of states was considered by several recent papers. In (Reijngoud and Endriss, 2012; van Ditmarsch et al., 2013) the assumption is that voters information sets can be described as a partition Π , as in the Aumann model. Reijngoud and Endriss say that a voter *has an incentive to Π -manipulate using ballot P'_i* (under profile \mathbf{P}), if she weakly gains by voting P'_i in every state that is “equivalent” to \mathbf{P} according to Π .²⁰ In the special case of Plurality, the definition coincides local dominance: Consider Def. 2, where we set S to be all states equivalent to \mathbf{a} under Π . Then a'_i S -dominates a_i iff i has an incentive to Π -manipulate using ballot a'_i . In the terminology of (van Ditmarsch et al., 2013), voter i *knows ‘de re’ that she can weakly successfully manipulate*. Our definition of local dominance also coincides with the definition of dominance in (Conitzer et al., 2011), which do not make any assumption on the “information set” S .

In our work the accessibility relation is defined by a distance metric and is *not* a partition. Still, many of the definitions in (Reijngoud and Endriss,

²⁰The definition is for arbitrary voting rules.

2012; van Ditmarsch et al., 2013) can be applied just the same in our case. In particular, a combination of these works can be used to extend the notion of local dominance to other voting rules.

7.3. Conclusion and future work

We see a unifying theory as the one we present as a productive step in the quest to understand voting. We hope that future researchers will find our theoretical framework useful for formulating new, more specific, voting behaviors. Furthermore, our particular distance-based model can serve as a strong baseline for competing theories. Experiments with human voters will be important to settle how close each of these theories comes in adequately describing human voting behavior.

On the technical level, we conjecture that stronger convergence properties can be proved; in particular, that there are no cycles in voting games with voters of the same type, and that a voting equilibrium exists even in games with heterogeneous voters.

We also believe that distance-based local dominance, with the necessary adaptations, can provide a useful non-probabilistic framework for uncertainty in other classes of games where there are natural distance metrics over states, such as congestion games.

Finally, insights based on our theory, for example on how voters' uncertainty level affects quality of the outcome, can be useful in designing better voting mechanisms.

Acknowledgments

We thank Joe Halpern, David Parkes, Yaron Singer, Greg Stoddard, and several anonymous reviewers for their valuable feedback.

Appendix A. Proofs

LEMMA 3. *Each of the metrics δ from $(\ell_1, \ell_\infty, \text{multiplicative})$ induces a function $\beta = \beta_{\delta, r} : \mathbb{N} \rightarrow \mathbb{N}$, where*

- *For every $\mathbf{a}, i \in N$, if $a_i \neq f_i^*$, then $W_i(\mathbf{a}, r) = \{c : s_{\mathbf{a}_{-i}}(c) \geq \bar{q} \beta(s(f_i^*))\}$. That is, possible candidates are all those whose score is above the threshold, which is a function of the score of the winner.²¹*

²¹We break ties with $\beta(s(f_i^*))$ as if we break ties with f_i^* .

- $\beta(s)$ is weakly increasing in s .

In particular, for δ_{ℓ_1} , $W_i(\mathbf{a}, r) = \overline{H}_{r+1}(\mathbf{s}_a)$ for any i s.t. $a_i \notin \overline{H}_{r+1}(\mathbf{s}_a)$. Similarly, for δ_{ℓ_∞} , $W_i(\mathbf{a}, r) = \overline{H}_{2r+1}(\mathbf{s}_a)$ for any i s.t. $a_i \notin H_{2r+1}(\mathbf{s}_a)$. A similar result can be proved for EMD and other ℓ_d norms, but the threshold would depend on the score of all candidates and not just the winner.

Proof. Consider first the ℓ_1 norm. We set $\beta(s) = \beta_{\ell_1}(s) = s - r - 1$. Clearly, f_i^* is a possible winner and $s(f_i^*) > \beta(\mathbf{a})$. Assume $c \neq f_i^*$, then $s(c) \neq \overline{Q} \beta(s(f_i^*))$. If $s(c) > \beta(s(f_i^*))$, consider the state $\mathbf{s}' \in S_i(\mathbf{a}, r)$ where c has r additional votes. We have that

$$s'(c) + 1 = s(c) + r + 1 > \beta(s(f_i^*)) + r + 1 = s(f_i^*),$$

thus $f(\mathbf{s}', c) = c$. In contrast, if $s(c) < \beta(s(f_i^*))$, then in any $\mathbf{s}' \in S_i(\mathbf{a}, r)$, $s'(f_i^*) - s'(c) > 1$ and thus c cannot win. Finally,

$$W_i(\mathbf{a}, r) = \{c : s(c) \geq \beta_{\ell_1}(s(f_i^*))\} = \{c : s(c) \geq s(f_i^*) - (r + 1)\} = \overline{H}_{r+1}(\mathbf{s}).$$

Similarly, for ℓ_∞ we set $\beta(s) = \beta_{\ell_\infty}(s) = s - 2r - 1$. The critical state \mathbf{s}' is where c gets r additional votes, and we subtract r votes from all other candidates (including f_i^*).

For the multiplicative distance, we set $\beta(s) = \lceil \lceil \frac{s}{1+r} \rceil / (1+r) \rceil - 1$. In the critical state we multiply the score of c by $(1+r)$, getting $s'(c) = \lfloor s(c)(1+r) \rfloor$, and divide the score of all other candidates by $(1+r)$, so e.g. for f_i^* , we get $s'(f_i^*) = \left\lceil \frac{s(f_i^*)}{1+r} \right\rceil$. Thus

$$\begin{aligned} f(\mathbf{s}', c) &= c && \iff \\ s'(c) + 1 &> s'(f_i^*) && \iff \\ \lfloor s(c)(1+r) \rfloor + 1 &> \left\lceil \frac{s(f_i^*)}{1+r} \right\rceil && \iff \\ s(c) &> \left\lceil \left\lceil \frac{s(f_i^*)}{1+r} \right\rceil / (1+r) \right\rceil - 1 = \beta(s(f_i^*)). \end{aligned}$$

□

Theorem 10. *Suppose that all voters are of type r and using any norm for which Lemmas 2,3 apply. Then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, for any singleton scheduler, they will always converge to an equilibrium in at most $n(m-1)$ steps.*

Proof. If the truthful state \mathbf{q} is stable, then we are done. Thus assume it is not. Let \mathbf{a}^t (and \mathbf{s}^t) be the voting profile after t steps from the initial truthful vote $\mathbf{a}^0 = \mathbf{q}$. Let $a_i \xrightarrow{i} a'_i$ be a move of voter i at state $\mathbf{s} = \mathbf{s}_{\mathbf{a}^t}$ to state $\mathbf{s}' = \mathbf{s}_{\mathbf{a}^{t+1}}$.

We claim that the following hold throughout the game.

1. $a_i \notin W_i(\mathbf{a}, r)$. Voters only leave non-possible winners.
2. After a step $a_i \xrightarrow{i} a'_i$ at time t , $a_i \notin W_j(\mathbf{a}^{t'}, r)$ at any later time t' , for any voter j .
3. $a'_i \prec_i a_i$. I.e., voters always “compromise” by voting for a less preferred candidate.
4. $\max_{a \in A} s'(a) \geq \max_{a \in A} s(a)$. I.e., the score of the winner never decreases.
5. For all j , all $t' > t$ after the first step of j , $W_j(\mathbf{a}^{t'}, r) \subseteq W_j(\mathbf{a}^t, r)$. I.e., the set of possible winners can only shrink (after the first move).

We prove this by a complete induction.

1. If this is the first move of i then $a_i = q_i$. Otherwise, by Lemma 2, a_i is the most preferred candidate of i in $W_i(\mathbf{a}^{t'}, r)$ where t' is the time when i last moved. By induction on (5), $W_i(\mathbf{a}, r) \subseteq W_i(\mathbf{a}^{t'}, r)$. So either a_i is not a possible winner in \mathbf{a} (and then we are done), or it must be the most preferred candidate in $W_i(\mathbf{a}, r)$. Assume, toward a contradiction, that $a_i \in W_i(\mathbf{a}, r)$, then there is a state in $S_i(\mathbf{a}, r)$ where i is pivotal between a_i and $f(\mathbf{a})$. With any other action, $f(\mathbf{a}) = c \prec_i a_i$ would win. Therefore a_i $S_i(\mathbf{a}, r)$ -beats any other candidate including a'_i . In particular, a'_i does not $S_i(\mathbf{a}, r)$ -dominate a_i , which is a contradiction.
2. The scores by which different voters determine possible winners are almost identical, and the score of a_i may differ by at most 1 vote between \mathbf{a}_{-i} and \mathbf{a}_{-j} . When i moves then by (1) and Lemma 3 $s_{\mathbf{a}_{-i}}(a_i) < \beta(s(f(\mathbf{a})))$, and thus $s_{\mathbf{a}_{-j}}(a_i) < \beta(s(f(\mathbf{a}))) + 1$. Thus while j may still consider a_i as a possible winner *before* i moves, a_i is no longer a possible winner for j after i moves, as

$$s'_{\mathbf{a}_{-j}}(a_i) = s_{\mathbf{a}_{-j}}(a_i) - 1 < \beta(s(f(\mathbf{a}))).$$

3. If this is the first move of i then this is immediate. Otherwise, by induction on Lemma 2 and (5), if $a'_i \succ_i a_i$, then i would prefer to vote for a'_i in his previous move, rather than to a_i .
4. As in (1), if i votes for $a_i = f(\mathbf{a}) \in \overline{H}_{r+1}(\mathbf{a})$, then a_i is i 's most preferred possible winner. Thus it cannot be locally dominated by any other candidate.

5. Since by (3) the score of the winner never decreases, the only way to expand $W_j(\mathbf{a}, r)$ is if some voter i added a vote to a candidate not in $W_j(\mathbf{a}, r)$. Recall that by Lemma 2, i only votes to candidates in $W_i(\mathbf{a}, r)$. Since $W_j(\mathbf{a}, r), W_i(\mathbf{a}, r)$ differ only by the votes of i and j , the only candidate in $W_i(\mathbf{a}, r) \setminus W_j(\mathbf{a}, r)$ can be a_j (the current vote of j).

However if j already moved once then a_j was a possible winner for j . Since at time t $a_j \notin W_j(\mathbf{a}, r)$, then some voter must have deserted a_j before time t . Then by (2) no voter would consider a_j a possible winner after time t , and by Lemma 2, no voter would vote for it before t' .

Note that if j have never moved, then it is possible that q_j is not a possible winner for j , but still gets a vote later from i .

Finally, by property (2), each voter moves at most $m - 1$ times before the game converges. \square

LEMMA 5. *Under the conditions of Theorem 4, either \mathbf{q} is stable, or in every state \mathbf{s}^t we have $|\overline{H}_{r+1}(\mathbf{s}^t)| > 1$. Also, in the last state \mathbf{a} either $|\overline{H}_r(\mathbf{s}_\mathbf{a})| = 1$, or all voters vote for possible winners. Any voter voting for $c \notin \overline{H}_{r+1}(\mathbf{s}^t)$ prefers $f(\mathbf{a})$ over any other candidate in $\overline{H}_{r+1}(\mathbf{s}^t)$.*

Proof. Note first that once $|\overline{H}_{r+1}(\mathbf{s})| = 1$, there are no strategic moves, as no candidate can challenge the winner. Thus a violation can occur only in the last step. Assume, toward a contradiction, that in the last step $a'_i = f(\mathbf{a})$, and $|\overline{H}_r(\mathbf{s})| = 1$ (but $|\overline{H}_{r+1}(\mathbf{s})| > 1$). However, since by (1) $a_i \notin \overline{H}_{r+1}(\mathbf{s})$, each candidate in $\overline{H}_{r+1}(\mathbf{s}) \setminus \{f(\mathbf{a})\}$ has the same score with and without i , and this score is at most $s(f(\mathbf{a})) - (r + 1)$. Therefore $f(\mathbf{a})$ also wins in any state in $S_i(\mathbf{a}, r)$ (no candidate is a threat to the winner). This means that $f(\mathbf{a})$ does not $S_i(\mathbf{a}, r)$ -beat a_i , in contradiction to a strategic move where $a'_i = f(\mathbf{a})$.

Finally, suppose that there is more than one candidate in $\overline{H}_r(\mathbf{s}_\mathbf{a})$. Then any voter i not voting for a possible winner sees himself potentially pivotal between the winner and the runner-up (there is a possible state where the runner-up wins if i keeps voting for a_i). Since i always strictly prefer on of them, this candidate will locally dominate a_i .

We can further see that if $|\overline{H}_r(\mathbf{s}_\mathbf{a})| = 1, |\overline{H}_{r+1}(\mathbf{s}_\mathbf{a})| > 1$, then any voter not voting for $\overline{H}_{r+1}(\mathbf{s})$ in equilibrium must favor the winner. Otherwise he would be potentially pivotal between the current winner and a better possible winner, and thus his most preferred possible winner would locally dominate his current vote. \square

PROPOSITION 6. *Suppose that all voters are of type r . Consider any group scheduler such that (1) any voter has some chance of playing as a singleton (i.e. this will occur eventually); (2) The scheduler always selects (an arbitrary subset of) voters with type 2 moves, if such exist. Then convergence is guaranteed from any initial state after at most $O(nm)$ singleton steps occur.*

Proof. Denote $N_w(\mathbf{a}) = \bigcup_{a \in \overline{H}_w(\mathbf{a})} \{i : a_i = a\}$ i.e. all voters voting for the top w candidates.

Note that there can be at most $n(m-1)$ consequent moves of type 2, regardless of the scheduler. Let \mathbf{a}^{t^*} be the state first reached when no voter has a type 2 step. If the state \mathbf{a}^{t^*} is stable, then we are done. Thus assume it is not.

Consider a group N' that moves at time t . By property (2) of the scheduler, either all of N' has type 1 moves, or all of it has type 2 moves. Thus we can classify all (group) steps to type 1 and type 2.

We define a *chunk potential function* α , where $\alpha(\mathbf{a}) = -n \times |\overline{H}_{r+1}(\mathbf{a})| + |N_{r+1}(\mathbf{a})|$. That is, the potential increases as the set of possible winners shrinks, but for a fixed size it increases with the total mass of voters for such candidates.

Consider a sequence of moves after time t^* . A *chunk* is composed of a type 1 step and all type-2 steps that follow until the next type 1 step (i.e. a type 1 step and then zero or more type 2 steps). We claim that (a) after every chunk α (as well as the score of the winner) does not decrease; (b) after a finite number of chunks α strictly increases.

We observe that in any type 1 move $\mathbf{a} \xrightarrow{N'} \mathbf{a}'$, $a'_i = \operatorname{argmin}\{Q_i(a) : a \in \overline{H}_{r+1}(\mathbf{s})\}$ for all $i \in N'$. We will prove by induction that starting from t^* , after every chunk $\mathbf{a} \xrightarrow{N'} \mathbf{a}' \xrightarrow{N''} \dots \xrightarrow{\hat{N}} \hat{\mathbf{a}}$:

1. $\max_{a \in M} \hat{s}(a) \geq \max_{a \in M} s(a)$. I.e., the score of the winner never decreases.
2. $\alpha(\hat{\mathbf{a}}) \geq \alpha(\mathbf{a})$.

We prove this by a complete induction.

Consider the type 1 move of the chunk. All voters vote for less preferred candidates, thus for all $i \in N'$, $s(a'_i) \geq s(a_i)$. Moreover, a voter only moves if $a_i \notin \overline{H}_r(\mathbf{a})$, since otherwise by he is potentially pivotal (as we show in the proof of Th. 4). Case I: for any $i \in N'$, there is no j s.t. $a'_j = a_i$. In this case all moves are essentially independent, and every i s.t. $a_i \notin \overline{H}_{r+1}(\mathbf{a})$ increases is added to $N_{r+1}(\mathbf{a}')$, thereby increasing α by 1. Every voter i s.t. $a_i \in \overline{H}_{r+1}(\mathbf{a})$ increases α by n , since $a_i \notin \overline{H}_{r+1}(\mathbf{a}')$. Thus in case I the chunk

potential strictly increases, and there are no followup type 2 steps (the chunk ends).

The problematic case is where there is $I \subseteq N'$, s.t. for all $i \in I$ there is $j \in N'$ with $a'_j = a_i$ (see Figure A.7). If it also holds that $a_i \in H_{r+1}(\mathbf{a})$, then after the move we have $a_i \in \overline{H}_{r+1}(\mathbf{a}')$. Then we have that while a_i was not a possible winner for i before the step, it is a possible winner for i after the step. Note however that only voters who moved can have a new possible winner, and it can only be the candidate they deserted. It is then possible that voter i now has a type 2 move. Thus some possibly empty subset $I' \subseteq I$ have type 2 moves, which is to return to their original candidates a_i (deserting other candidates in $\overline{H}_{r+1}(\mathbf{a}')$). After the first type 2 move there may other voters $I'' \subseteq I \setminus I'$ that want to return and so on. However no other voter has a new type 2 move since all of $N' \setminus I$ vote for their most preferred possible winner in \mathbf{a}' ; and any $i' \in N \setminus N'$ does not have a type 2 move since they did not have one in \mathbf{a} and there are no new possible winners.

So every further step 2 in the chunk rolls back some of the first type 1 steps. At the end of the chunk we are left with a set $N'' = N' \setminus (I' \cup I'' \cup \dots)$, where N'' performed a type 1 step and all other voters vote as in \mathbf{a} . If $N'' \neq \emptyset$ then by the previous case α strictly increases (and the score of the winner does not decrease). Clearly if $N'' = \emptyset$ then $\hat{\mathbf{a}} = \mathbf{a}$ and thus α does not change. However note that the type 1 step is a set of disjoint cycles. For N'' to be empty, each of these cycles must be contained completely in I' or I'' , etc.: a voter that does not roll back his action $a_j \xrightarrow{j} a'_j$ at the same time with the voter who joined a_j , will not be able to roll back at a later type 2 step, since a_j will no longer be a possible winner for j . Thus a singleton type 2 move (which cannot be a cycle) means that $N'' \neq \emptyset$.

Since eventually there will be singleton move (either type 1 or type 2), the same cycle repeat forever, and α must increase. Clearly it cannot increase more than nm times. \square

PROPOSITION 7. *Suppose that each voter i is of type $L(r, k_i)$ or $T(r, k_i)$, where $k_i > r$. Then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, they will always converge to an equilibrium in at most $3nm$ steps.*

Proof. We first prove for truth-biased voters. Consider a $T(r, k_i)$ voter i . A truth-bias move can only occur when i has no strategic moves. By Lemma 5, $|\overline{H}_{r+1}(\mathbf{s})| > 1$, and by Lemma 3 this means there are at least two possible winners for i . If some of them locally dominate a_i , then i would have a strategic move. Thus if i makes a truth-bias move he is in one of two situations:




Figure A.7: An example of a group scheduler, with $r = 1$. Voters 1,2,3 rank d last, and rank their current vote first. In the initial state $\mathbf{s} = (1, 1, 1, 3)$, we have $H_{r+1}(\mathbf{s}) = \{a, b, c\}$. Step 1 is a compromise move, where $N' = \{1, 2, 3\}$. After this move, a, b and c are all still possible winners. Now all three voters have an opportunity move, which is to go back to their previous vote. If the scheduler keeps selecting $\{1, 2, 3\}$ then we would have a cycle going forever. In this example the scheduler picks $I' = \{1\}$ (step 2), which ends the chunk. Now there is only one voter with a strategic move (voter 2), so the next chunk has just one step. After step 3 no voter has a strategic move, so the game is in equilibrium.

(a) when i is already voting for his most preferred candidate in $\overline{H}_{r+1}(\mathbf{s})$; (b) $a_i \notin \overline{H}_{r+1}(\mathbf{s})$, but none of the possible winners $S_i(\mathbf{a}, r)$ -dominates a_i . Denote such moves by type-a and type-b, respectively.

We first argue that there are no type-a truth-bias moves. Indeed, we have $a_i \in \overline{H}_{r+1}(\mathbf{s}) \subseteq \overline{H}_{k_i}(\mathbf{s})$. If $a_i \neq f(\mathbf{a})$, there is a state $\mathbf{s}^* \in S_i(\mathbf{a}, k_i)$ where a_i wins if getting the vote of i , and otherwise $f(\mathbf{a})$ wins. If $a_i = f(\mathbf{a})$, then consider some other $b \in \overline{H}_{r+1}(\mathbf{s})$ (such a b exists according to Lemma 5). Then there is a state $\mathbf{s}^* \in S_i(\mathbf{a}, k_i)$ where $a_i = f(\mathbf{a})$ wins if getting the vote of i , and otherwise b wins. In either case, a_i $S_i(\mathbf{a}, k_i)$ -beats q_i , and thus there is no type-a move. Note that this is where we apply the assumption that $k_i > r$.

Type-b moves are possible, and thus we need to show that invariants (3) and (4) in the proof of Theorem 4 still hold. I.e., that a truth-bias move cannot cause the winner to lose votes, and cannot expand the possible winners set. Clearly the winner cannot lose score, since $a_i \notin \overline{H}_{r+1}(\mathbf{a})$. It is left to show that $a'_i = q_i$ cannot become a possible winner. Let $N'_{q_i} \subseteq N$ be all core supporters of q_i that vote strategically in \mathbf{a} . For all $j \in N'_{q_i}$ (including i), when j deserted q_i then q_i was no longer a possible winner for j (i.e., its score was at least $r+2$ below the winner). Since then, the score of the winner does not decrease. Thus even if all of N'_{q_i} return to q_i , the gap between q_i and $f(\mathbf{a})$ would be at least $r+2$.

For the bound on the number of steps, denote by $R_i, K_i \in \mathbb{N}$ the number of strategic moves and Truth-bias moves of voter i , respectively. Observe that a strategic move can only occur after the set of possible winners shrinks—unless it is the first move or it follows a truth-bias move. Thus $R_i \leq (K_i + 1) + (m - 1) \leq K_i + m$. A truth-bias can only come after the set of possible winner shrinks as well, thus $K_i \leq m - 1$. In total, $R_i + K_i \leq 2K_i + m \leq 3m$. So all voters together can do at most $3nm$ moves.

We now turn to prove that convergence still holds when adding lazy-biased voters. Assume first that $|\overline{H}_{r+1}(\mathbf{q})| > 1$. Then the proof is essentially the same, replacing q_i with \perp for lazy-biased voters.

However the proof breaks when $|\overline{H}_{r+1}(\mathbf{q})| = 1$, as “lazy” supporters of the winner may abstain. Thus this case requires a special treatment. We will show that every lazy voter plays at most once, and then the game ends.

Let k^t be the gap between $a^* = f(\mathbf{q})$ and its closest runner-up (i.e., $k^t = \min\{k > 0 : H_k(\mathbf{s}^t) \neq \emptyset\}$). By our assumption $k^0 > r + 1$. At time t , voter $i = i_t$ will choose to abstain if and only if $k_i \leq k^t - 2$. To see why, if $k_i \geq k^t - 1$, then $a^* S_i(\mathbf{a}^t, k_i)$ -beats \perp (by removing k_i votes from a^* , i becomes pivotal between a^* and the runner-up). However, a voter

who already abstains, will choose to enter only if $k^t \leq r + 1$ (or $k^t \leq r$, for supporters of a^*). Thus k^t may either increase or decrease by 1 in every step. However in any case k^t cannot go below $r + 2$. We can show this by induction: the base case follows since $|\overline{H}_{r+1}(\mathbf{q})| = 1$. In any later step, $k^{t+1} < k^t$ if and only if a supporter i of a^* abstains, which occurs iff $r < k_i \leq k^t - 2$. Thus $k^t \geq r + 3$, and $k^{t+1} \geq r + 2$.

Finally, after each voter made his decision whether to abstain or not, there are no strategic moves, as $k^t \geq r + 2$ entails $|\overline{H}_{r+1}(\mathbf{a}^t)| = 1$. In this case, the game will converge after at most n moves. \square

Appendix B. Simulations

Appendix B.1. Distributions of preference profiles

- **Uniform:** Also known as the *impartial culture* distribution, this is the simplest distribution to study. While people’s votes are rarely distributed at random, the uniform distribution allows more confidence that our results are not particular and specific to the distributions analyzed, and is thus often used in simulations of voting (Nurmi, 1992).
- **Single-peaked:** This distribution assigns each candidate a point on the interval $[0, 1]$, and each voter is randomly assigned a point on the interval, which defines its preferences — it prefers candidates closer to its point. This distribution has been long used in sociological and political research (as resembling the common right-left political axis) (Kedar, 2014), but has also been widely examined in game theoretic scenarios. A particular interesting property is that for single-peaked preferences can be aggregated using strategy-proof mechanisms. The most prominent such mechanism is the *median vote* (Sprumont, 1991).²²
- **Polya-Eggenberger urn:** This model was developed and used to model the grouping of much of society to major homogenous groups (Berg, 1985; Walsh, 2010; Rothe and Schend, 2012). In a k -urn model, k preference orders are chosen, and an urn is built to let voters choose preference orders from it. Each of the k chosen preferences gets $\frac{1}{k+1}$ of the preference orders in the urn, with remaining $\frac{1}{k+1}$ of the urn filled

²²We also tried simulations with *single-dipped* preferences, where the voter’s most preferred candidates are at the extreme. However, in such profiles the truthful vote has only two candidates with positive support (the extremes), and no voter ever has an incentive to move.

by all non-selected preference orders. Preferences chosen from this urn have significant likelihood to be of the k main groups. In this work, we used the 2-urn and 3-urn model.

- **Riffle:** In a riffle model we get preferences of each voter by interleaving two separate preference orders on subsets of candidates in an independent manner. Huang and Guestrin (2009) showed real-world elections which resembled this distribution.
- **Plackett-Luce:** In the Plackett-Luce model each candidate has an intrinsic cardinal value in the interval $[0, 1]$ (the “ground truth”). Each vote is then sampled from a particular distribution which adds noise to the true ranking (Soufiani et al., 2012).

German election data. We used data from three polls based on the German National Election Study, from 1969, 1972, and 1976.²³ The German election datasets had three candidates: The Christian Democratic Union (CDU), the Social Democratic Party (SDP), and the Free Democratic Party (FDP); we used $n = 100$ voters whose preferences (from the six possible orders) were taken from (Regenwetter et al., 2006).

PrefLib. We used all 225 complete preference profiles available from PrefLib.org. Most instances have 3 or 4 candidates (over 10 in some), and several hundred voters.

Appendix B.2. Methods

We simulated voting in an iterative setting, where voters start from a particular state, and then iteratively make strategic moves until convergence. We constructed a simulator that enables us to control the following features of the simulation. First, determine the parameters of the preference profile:

- Number of voters. We used $n \in \{10, 20, 50\}$ (also $n = 100$ in a few simulations).
- Number of candidates. We used $m \in \{3, \dots, 8\}$.
- Distribution of preferences. We used all the 6 distributions described above.

²³Since these were election to the parliament and not a single-winner elections, it is impossible to compare the simulation results to the actual outcome.

Ignoring the $n = 100$ simulations, this defines $3 \times 6 \times 6 = 108$ distinct distributions. From each such distribution we generated 200 preference profiles, to a total of 21,600 profiles. Then, we determine the parameters of the strategic model:

- Distance metric used for accessibility relation. We used the ℓ_1 norm and the multiplicative distance.
- The uncertainty parameter r which determines the radius of local dominance. For the ℓ_1 distance, we varied r in $\{0, 1, \dots, 15\}$; for multiplicative we varied r in $\{\frac{1}{n}, \dots, \frac{4}{n}, 0.3, 0.5, 0.7\}$. We also used $r = m$ as a baseline value, where there is no strategic behavior (the outcome is the truthful Plurality outcome).
- For voters with truth-bias or lazy-bias, the parameter k was varied as well from $r + 1$ (or $r + \frac{1}{n}$ for multiplicative) to $2r$.
- The simulator allows the creation of voters of different types. However in most of our simulations all voters had the same type (but different preferences).

Finally, there are setting of the simulation itself:

- The initial voting profile. In most simulations this was the voters' truthful profile.
- The scheduler. In most simulations we used a singleton scheduler, which selects the next voter uniformly at random. We also used a group scheduler, which randomly selects a subset of at most $n/2$ voters and lets all of them make independent strategic moves (if they have such).

Basic simulations. For each profile, we ran simulations with strategic voters using the ℓ_1 metric, ranging the value of r . For each combination of preference profiles and voters' type, we repeated the simulation 100 times with a random singleton scheduler, recording all the equilibria that were attained from this profile. In this setting, we know by Theorem 4 that every simulation must converge. For profiles with $m \leq 5$, we conducted the same procedure with the multiplicative metric as well.

Concurrent voting. We repeated a subset of the simulations above with a group scheduler. We used all the above distributions, with $n = 50, m = 5$, and all values of r . We then compared the results to the corresponding results under a singleton scheduler.

Random initial state. We repeated our simulations where the initial state is sampled uniformly at random from all voting profiles (regardless of the preference profile). Thus for each preference profile we ran 100 simulations, each with a different initial state and a different scheduler.

Diverse types. Since actual societies are likely to contain voters of different types, we repeated the same simulations with a diverse population. In these simulations the value r_i for each voter was sampled uniformly from $\{0, 1, \dots, n/m\}$.

Note that in the three batches of simulations above, we had no formal guarantee that the game converges. However, in practice all simulations converged to an equilibrium.

Truth-biased and lazy-biased societies. For $m \leq 5$ we repeated the simulations while varying the values of k in addition to r . In that process we simulated truth-biased and lazy-biased societies.

Non-truthful starting profiles. For $n = 20, m = 4$, we examined the possibility of changing the the starting position of players to a random one. A different starting profile was assigned every run of a profile (to enable examination how many different winners are reached using our dynamic).

Observed variables. For every generated profile, we measured the following variables (all averaged over 100 simulations with a random singleton scheduler). We then averaged again over all 200 generated profiles of a given distribution.

NumStep The number of steps from the initial (truthful) profile to convergence.

NumStates The number of distinct equilibrium outcomes (for the same preference profile), in terms of voting profiles.

NumWinners The number of distinct equilibrium outcomes (for the same preference profile), in terms of winner's identity.

WinnerConsistency The maximal fraction of simulations (for the same preference profile) that ended with the same winner.

PluralityAgreement The fraction of simulations where the winner was the original Plurality winner.

BordaAgreement The fraction of simulations where the winner was the Borda winner.

CoplandAgreement The fraction of simulations where the winner was the Copland winner.

MaximinAgreement The fraction of simulations where the winner was the Maximin winner.

CondorcetAgreement The fraction of simulations where the winner was the Condorcet winner, when one exists (not counted otherwise).

SocialWelfare The relative rank of the winner, according to its borda score (lower is better). Equivalently: the complement of the average social welfare of voters, assuming Borda utilities.

Gap1-2 The ratio between the score of the winner and the score of the runner-up ($s(c_1)/s(c_2)$, where $c_i = \operatorname{argmax}_{c \neq c_j, j < i} s(c)$.)

Gap2-3 The ratio between the score of the second and the third candidates ($s(c_2)/s(c_3)$).

TotalDuverger Fraction of simulations where at most two candidates received votes (1 iff $s(c_3) = 0$).

RelativeDuverger The fraction of votes for the two leading candidates $((s(c_1) + s(c_2))/n)$.

For the Plackett-Luce distribution, we also measured *WinnerGroundRank*, which is the rank of the winner according to the ground truth used to generate the profile (between 0 and $m - 1$).




Figure B.8: Ratio of games where the winner was the same winner as plurality for Riffle distribution simulation (*PluralityAgreement*)




Figure B.9: Improvement in the objective quality of the winner (*Winner-GroundRank*), according to the ground truth. Note that when there is some strategic interaction but not as much as in peak r , winner quality might decrease.




Figure B.10: Same as Figure B.9, but with random initial state. We can see strategic behavior (around peak r) yields better outcome on average than truthful voting, even when starting from a random state.




Figure B.11: The fraction of simulations in the single-peaked distribution, where the winner is the median candidate (*Condorcet Agreement*), relatively to the truthful baseline.

References

- Abramson, P. R., Aldrich, J. H., Paolino, P., Rohde, D. W., 1992. “sophisticated” voting in the 1988 presidential primaries. *American Political Science Review* 86 (01), 55–69.
- Aldrich, J. H., 1993. Rational choice and turnout. *American Journal of Political Science*, 246–278.
- Alvarez, R. M., Nagler, J., 2000. A new approach for modelling strategic voting in multiparty elections. *British Journal of Political Science* 30 (1), 57–75.
- Apt, K. R., Simon, S., 2012. A classification of weakly acyclic games. In: *SAGT’12*. pp. 1–12.
- Aumann, R. J., 1995. Backward induction and common knowledge of rationality. *Games and Economic Behavior* 8 (1), 6–19.
- Aumann, R. J., 1999. Interactive epistemology i: knowledge. *International Journal of Game Theory* 28 (3), 263–300.
- Berg, S., 1985. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice* 47 (2), 377–387.
- Bicchieri, C., Antonelli, G. A., 1995. Game-theoretic axioms for local rationality and bounded knowledge. *Journal of Logic, Language, and Information* 4 (2), pp. 145–167.
- Brock, W. A., Durlauf, S. N., 2001. Discrete choice with social interactions. *The Review of Economic Studies* 68 (2), 235–260.
- Calvert, R. L., 1985. Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, 69–95.
- Chopra, S., Pacuit, E., Parikh, R., 2004. Knowledge-theoretic properties of strategic voting. Presented in *JELIA-04*, Lisbon, Portugal.
- Conitzer, V., Walsh, T., Xia, L., 2011. Dominating manipulations in voting with partial information. In: *AAAI*. Vol. 11. pp. 638–643.
- Dekel, E., Piccione, M., 2000. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy* 108 (1), 34–55.
- Desmedt, Y., Elkind, E., 2010. Equilibria of plurality voting with abstentions. In: *ACM-EC’10*. pp. 347–356.

- Dhillon, A., Lockwood, B., 2004. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior* 46, 55–75.
- Downs, A., 1957. *An economic theory of democracy*.
- Dutta, B., Sen, A., 2012. Nash implementation with partially honest individuals. *Games and Economic Behavior* 74 (1), 154–169.
- Duverger, M., 1954. *Political Parties: Their Organization and Activity in the Modern State*. New York: John Wiley. Y.
- Edlin, A. S., Gelman, A., Kaplan, N., 2007. Voting as a rational choice: Why and how people vote to improve the well-being of others. *Rationality and society* 1.
- Falik, D., Meir, R., Tennenholtz, M., 2012. On coalitions and stable winners in plurality. In: *WINE’12*. pp. 256–269.
- Farquharson, R., 1969. *Theory of Voting*. Yale Uni. Press.
- Feddersen, T. J., Sened, I., Wright, S. G., 1990. Rational voting and candidate entry under plurality rule. *American Journal of Political Science* 34 (4), 1005–1016.
- Felsenthal, D. S., Rapoport, A., Maoz, Z., 1988. Tacit co-operation in three-alternative non-cooperative voting games: a new model of sophisticated behaviour under the plurality procedure. *Electoral Studies* 7 (2), 143–161.
- Ferejohn, J. A., Fiorina, M. P., 1974. The paradox of not voting: A decision theoretic analysis. *The American political science review*, 525–536.
- Fisher, S. D., 2004. Definition and measurement of tactical voting: the role of rational choice. *British Journal of Political Science* 34 (1), 152–166.
- Grandi, U., Loreggia, A., Rossi, F., Venable, K. B., Walsh, T., 2013. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In: *ADT’13*. pp. 181–192.
- Huang, J., Guestrin, C., December 2009. Riffled independence for ranked data. In: *NIPS’09*.
- Johnson, E. J., Camerer, C., Sen, S., Rymon, T., 2002. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory* 104 (1), 16–47.
- Kahneman, D., Knetsch, J. L., Thaler, R. H., 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives* 5 (1), 193–206.

- Kedar, O., January 2014. Voting for Policy, Not Parties. Cambridge Studies in Comparative Politics. Cambridge University Press.
- Laslier, J.-F., 2009. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics* 21 (1), 113–136.
- Manski, C. F., 1993. Identification of endogenous social effects: The reflection problem. *The review of economic studies* 60 (3), 531–542.
- McKelvey, R. D., Niemi, R., 1978. A multistage representation of sophisticated voting for binary procedures. *Journal of Economic Theory* 18, 1–22.
- Meir, R., Polukarov, M., Rosenschein, J. S., Jennings, N. R., 2010. Convergence to equilibria in plurality voting. In: AAAI’10.
- Merrill, S., 1982. Strategic voting in multicandidate elections under uncertainty and under risk. In: *Power, voting, and voting power*. Springer, pp. 179–187.
- Messner, M., Polborn, M. K., 2002. Robust political equilibria under plurality and runoff rule. Mimeo, Bocconi University.
- Myerson, R. B., Weber, R. J., 1993. A theory of voting equilibria. *The American Political Science Review* 87 (1), 102–114.
- Niemi, R. G., Frank, A. Q., 1982. Sophisticated voting under the plurality procedure. In: *Political Equilibrium*. Springer, pp. 151–172.
- Nurmi, H., 1992. An assessment of voting system simulations. *Public Choice* 73 (4), 459–487.
- Obraztsova, S., Markakis, E., Thompson, D. R., 2013. Plurality voting with truth-biased agents. In: SAGT’13. pp. 26–37.
- Owen, G., Grofman, B., 1984. To vote or not to vote: The paradox of nonvoting. *Public Choice* 42 (3), 311–325.
- Palfrey, T. R., Rosenthal, H., 1983. A strategic calculus of voting. *Public Choice* 41 (1), 7–53.
- Regenwetter, M., Grofman, B., Marley, A., Tsetlin, I., 2006. Behavioral social choice. Cambridge: Cambridge University Press.
- Reijngoud, A., Endriss, U., 2012. Voter response to iterated poll information. In: AAMAS’12. pp. 635–644.
- Reyhani, R., Wilson, M. C., Khazaei, J., 2012. Coordination via polling in plurality voting games under inertia. In: COMSOC’12.

- Riker, W. H., Ordeshook, P. C., 1968. A theory of the calculus of voting. *The American Political Science Review*, 25–42.
- Rothe, J., Schend, L., 2012. Control complexity in bucklin, fallback, and plurality voting: An experimental approach. In: *Experimental Algorithms*. Springer Verlag, pp. 356–368.
- Sertel, M. R., Sanver, M. R., 2004. Strong equilibrium outcomes of voting games are the generalized condorcet winners. *Social Choice and Welfare* 22 (2), 331–347.
- Silberman, J., Durden, G., 1975. The rational behavior theory of voter participation. *Public Choice* 23 (1), 101–108.
- Soufiani, H. A., Parkes, D. C., Xia, L., 2012. Preference elicitation for general random utility models. In: *NIPS’12*. pp. 126–134.
- Sprumont, Y., March 1991. The division problem with single-peaked preferences: A characterization of the uniform allocation rule. *Econometrica* 59 (2), 509–519.
- Thompson, D. R., Lev, O., Leyton-Brown, K., Rosenschein, J., 2013. Empirical analysis of plurality election equilibria. In: *AAMAS’13*. pp. 391–398.
- Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: Heuristics and biases. *science* 185 (4157), 1124–1131.
- van Ditmarsch, H., Lang, J., Saffidine, A., 2013. Strategic voting and the logic of knowledge. In: *TARK’13*. To appear.
- Walsh, T., 2010. An empirical study of the manipulability of single transferable voting. In: *ECAI’10*. pp. 257–262.




Figure B.12: The Difference between *RelativeDuverger* in the equilibrium outcome (the fraction of voters voting for the two leaders), and the truthful baseline.




Figure B.13: Number of steps to convergence (*NumStep*).